Connected computers are better! How' s it done?

Networking ...

Based on slides by Lawrence Snyder
University of Washington, Seattle

Networks…

Computers are useful alone, but are better when connected (networked)

- Access more information and software than is stored locally
- Help users to communicate, exchange information…changing ideas about social interaction
- Perform other services—printing, Web, email, texting, mobile, etc.

History of the Internet

by Melih Bilgil

1 year ago

http://www.vimeo.com/2696386

Video Summary

- timesharing vs batch
- TCP verifies file transfer
- packet switching
- decentralized network
- replaced radio communication
- packets pass through intermediate nodes unexamined
- x.25 allowed data over telephone networks
- OSI – layered model
- TCP/IP

Network Structure

Networks are structured differently based (mostly) on distance between computers:

- Local area network (LAN)
  - Small area: room or building
  - Either wired (Cu or fiber) or wireless
- Wide area networks (WAN)
  - Large area: more than 1 km
  - Fiber-optic, copper transmission lines, μ-wave, satellite
- Metropolitan area networks (MAN)
  - Neighborhood or several blocks of business district
  - Private service provider owns network

Protocol Rules!

To communicate computers need to know how to set up the info to be sent and interpret the info received

- Communication rules are a *protocol*
- Example protocols
  - EtherNet—for physical connection in a LAN
  - TCP/IP—for Internet—transmission control protocol / internet protocol
  - HTTP—for Web—hypertext transfer protocol

OSI 7-Layer Network Model

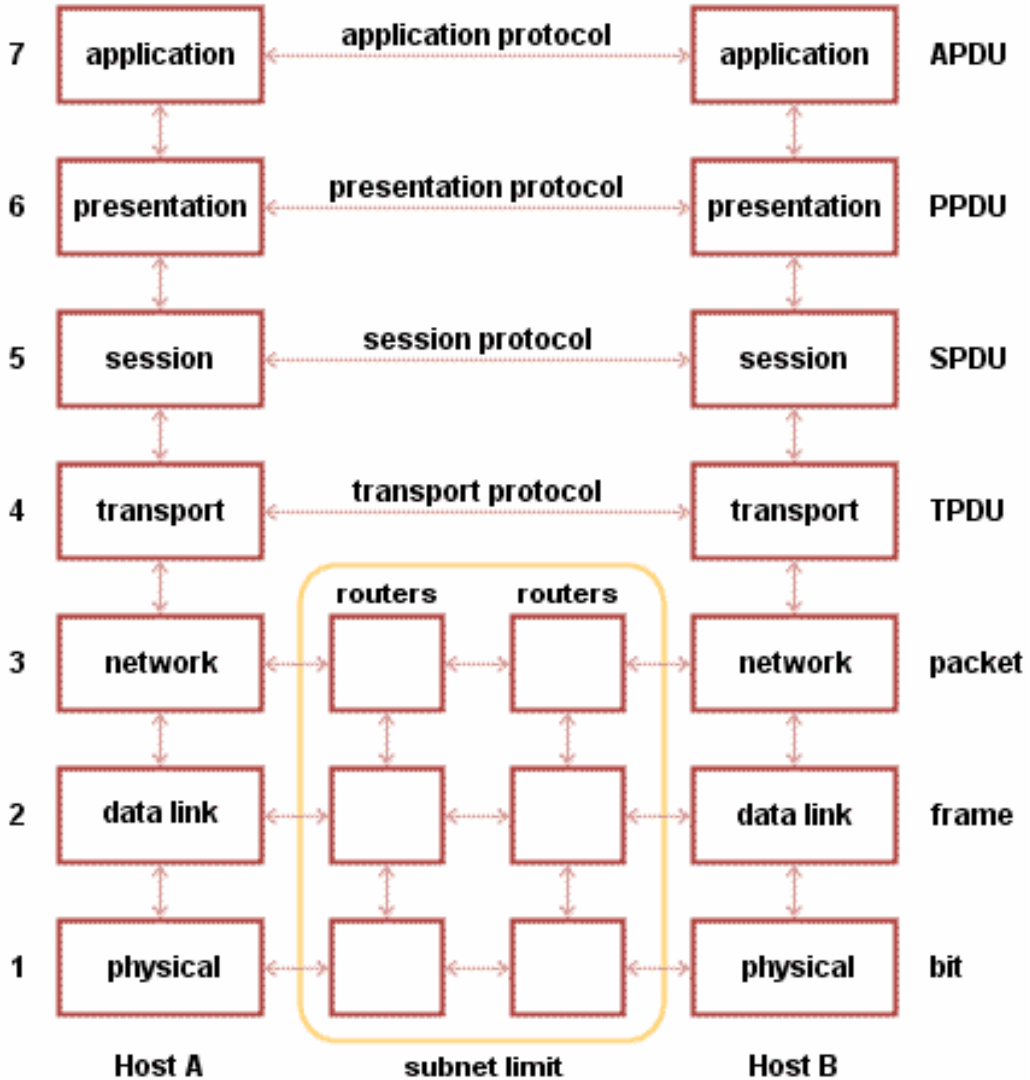

7. Application layer

NNTP · SIP · SSI · DNS · FTP · Gopher · HTTP · NFS · NTP · SMPP · SMTP · SNMP · Telnet · DHCP · Netconf · (more)

6. Presentation layer

MIME · XDR

5. Session layer

Named pipe · NetBIOS · SAP · PPTP · RTP · SOCKS · SPDY · TLS/SSL

4. Transport layer

TCP · UDP · SCTP · DCCP · SPX

3. Network layer

IP (IPv4 · IPv6) · ARP · ICMP · IPsec · IGMP · IPX · AppleTalk

2. Data link layer

ATM · SDLC · HDLC · CSLIP · SLIP · GFP · PLIP · IEEE 802.2 · LLC · L2TP · IEEE 802.3 · Frame Relay · ITU-T G.hn DLL · PPP · X.25

1. Physical layer

EIA/TIA-232 · EIA/TIA-449 · ITU-T V-Series · I.430 · I.431 · PDH · SONET/SDH · PON · OTN · DSL · IEEE 802.3 · IEEE 802.11 · IEEE 802.15 · IEEE 802.16 · IEEE 1394 ·

LAN in the Lab

EtherNet is a popular LAN protocol

- It uses a “party” protocol

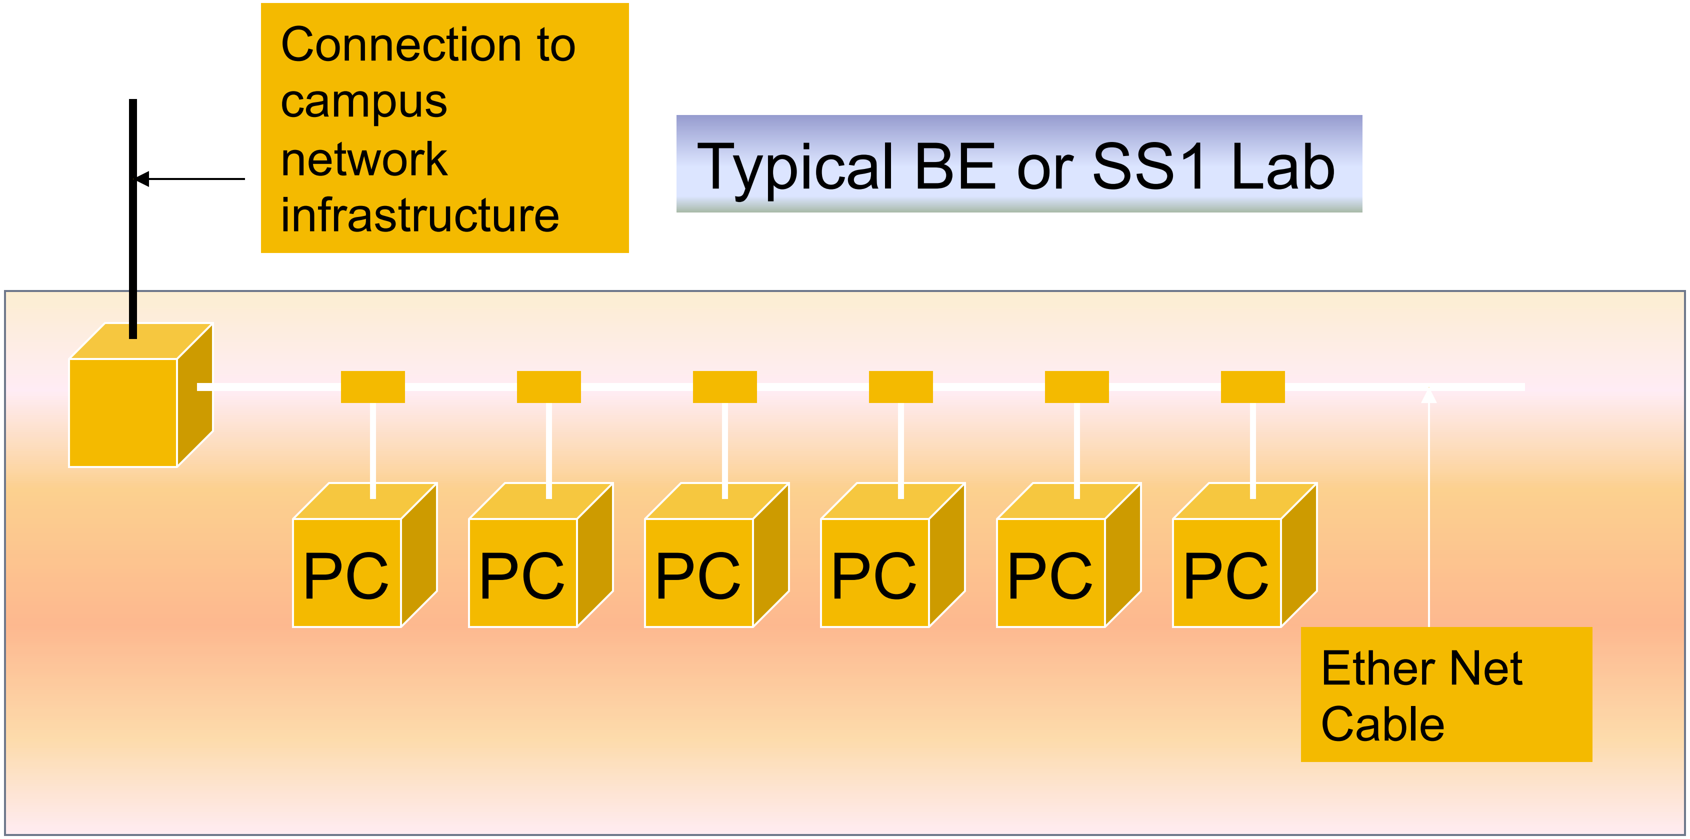

Campus & The World

The campus subnetworks interconnect computers of the UCSC domain which connects to Internet via a gateway

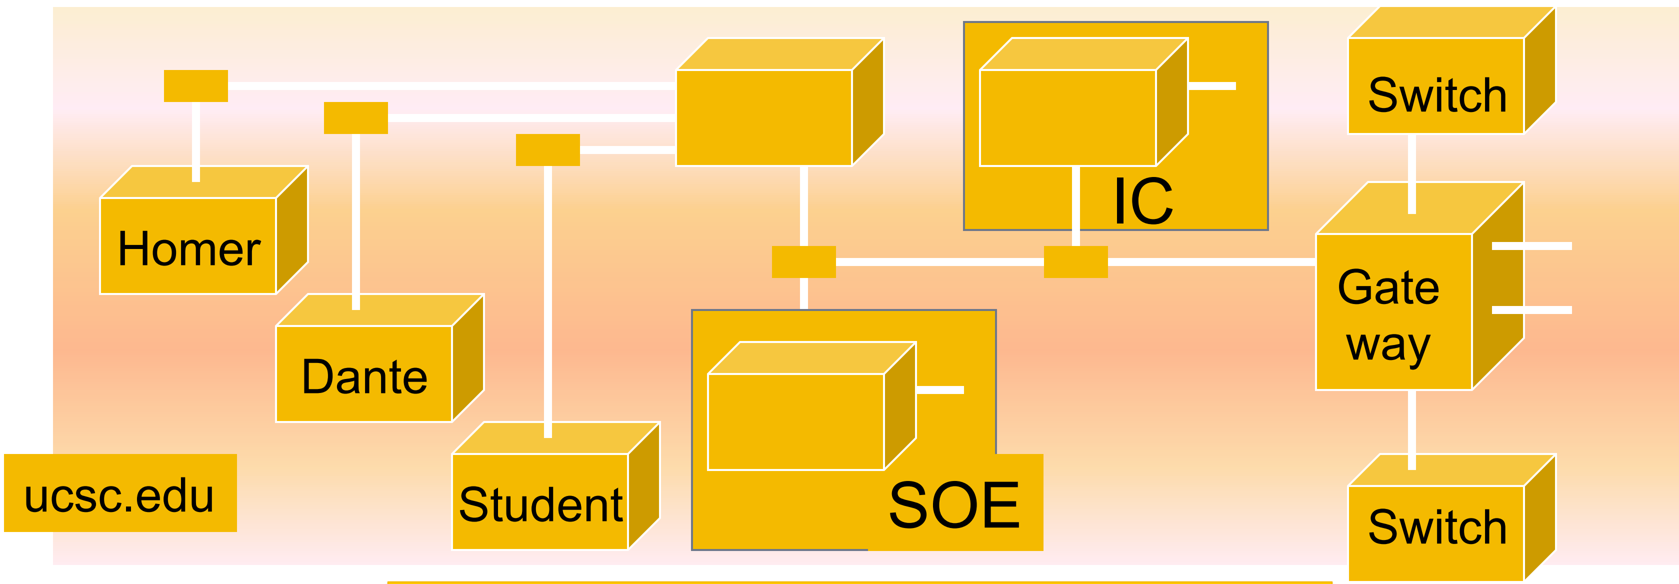

All communication by TCP/IP

IP—Like Using Postcards

Information is sent across the Internet using IP—Cerf uses postcard analogy

- Break message into fixed size units
- Form IP packets with destination address, sequence number and content

addr	#	data

- Each makes its way separately to destination, possibly taking different routes
- Reassembled at destination forming msg

Key Point: Taking separate routes lets packets by-pass congestion and out-of-service switches; packet reassembly discovers lost packets; ask for resend

http://www.youtube.com/watch?v=PBWhzz_Gn10

Video Summary

- IP Packets (following their journey)
- Proxy Server censors some URLs
- Router moves from one network to another
- Switch – similar in function to a router
- Firewall – more screening of IP packets at boundary of local network and the Internet
- IP Packets that get lost are resent
- Ping of death
- Ports

Routes in the Internet

A. Sending a big message across the internet is like a train with many cars (smaller packets) that are linked together to form a long train (the entire message).
B. Sending a big message across the internet is like sending a bunch of cars (smaller packets) that each take their own route to the destination and get reassembled after they all get there. Some might not even make it the first time.

Picture of Portion of I'net

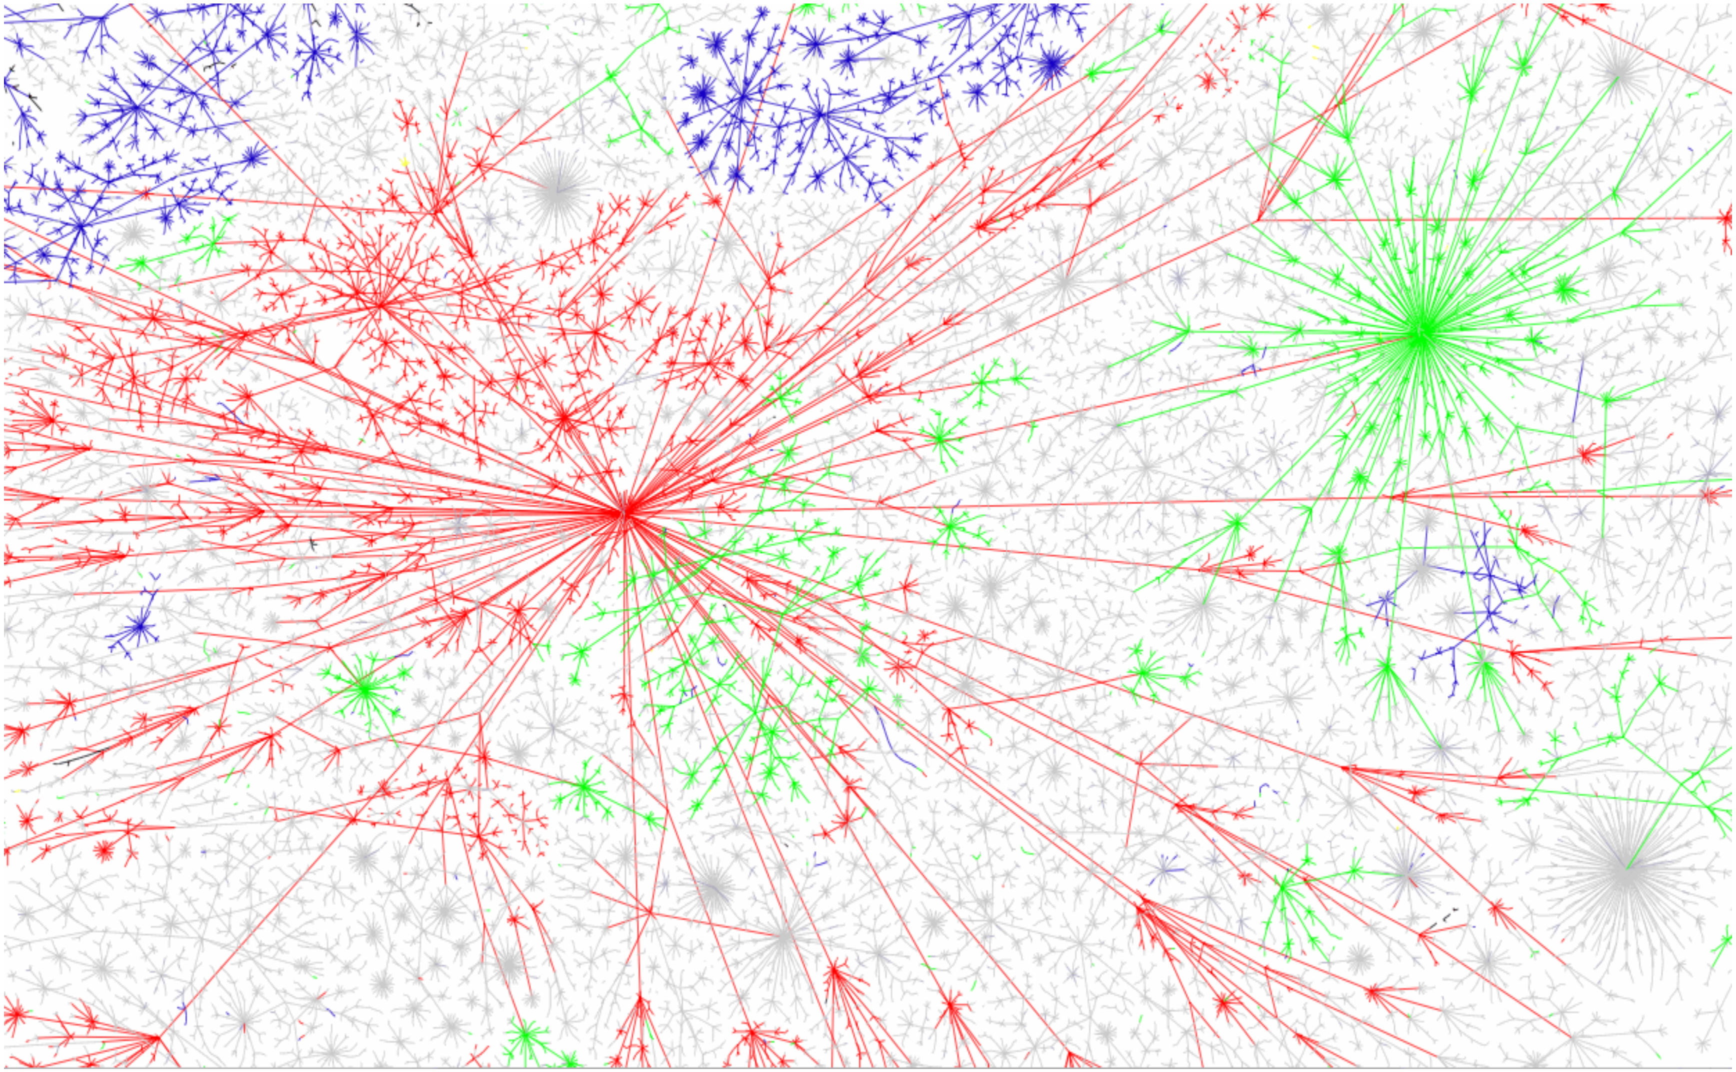

A Quick Trip to U. Florida

Table: Traceroute to cise.ufl.edu

Hop		IP Address	Node Name	Location		ms		Network
0		10.0.1.1	-			4		[Local Network]
1		10.0.1.1	-			3		[Local Network]
2								
3		68.87.205.1	-	Mt Laurel, usa		16		[Network for 68.87.205.1]
4		68.85.240.101	be-70-ar01.burien.wa.seattle.comcast.net	Mt Laurel, usa		13		[Network for 68.85.240.101]
5		68.85.240.69	be-30-ar01.seattle.wa.seattle.comcast.net	Seattle, WA, USA		13		[Network for 68.85.240.69]
6		68.86.90.213	pos-0-5-0-0-cr01.seattle.wa.ibone.comcast.net	Seattle, WA, USA		10		[Network for 68.86.90.213]
7		68.86.85.206	pos-0-8-0-0-cr01.portland.or.ibone.comcast.net	Portland, OR, USA		16		[Network for 68.86.85.206]
8		68.86.85.197	pos-1-15-0-0-cr01.sacramento.ca.ibone.comcast	Sacramento, CA, USA		26		[Network for 68.86.85.197]
9		68.86.85.181	pos-0-9-0-0-cr01.sanjose.ca.ibone.comcast.net	San Jose, CA, USA		34		[Network for 68.86.85.181]
10		154.54.11.105	te3-3.mpd01.sjc04.atlas.cogentco.com	San Jose, CA, USA		45		PSINet, Inc.
11		154.54.0.177	te9-1.ccr02.sfo01.atlas.cogentco.com	San Francisco, CA, USA		33		PSINet, Inc.
12		154.54.3.137	te3-8.ccr01.lax01.atlas.cogentco.com	Los Angeles, CA, USA		62		PSINet, Inc.
13		154.54.0.226	te3-8.ccr01.iah01.atlas.cogentco.com	Houston, TX, USA		97		PSINet, Inc.
14		154.54.24.194	te3-2.ccr01.mia01.atlas.cogentco.com	Miami, FL, USA		110		PSINet, Inc.
15		154.54.1.186	te3-3.ccr01.mia03.atlas.cogentco.com	Miami, FL, USA		114		PSINet, Inc.
16		38.112.31.66	florida_lambdarail_llc.demarc.cogentco.com	Washington, DC, USA		111		PSINet, Inc.
17		198.32.155.10	tpa-flrcore-7609-1-te21-1.net.flrnet.org	Marina del Rey, usa		124		EP.NET, LLC.
18		198.32.173.161	tlh-flrcore-7609-1-te41-1907.net.flrnet.org	Marina del Rey, usa		122		EP.NET, LLC.
19		198.32.173.162	ctx36-ewan-msfc-1-v1907-1.ns.ufl.edu	Marina del Rey, usa		203		EP.NET, LLC.
20		128.227.236.85	ctx36-nexus-msfc-1-v50-1.ns.ufl.edu	Gainesville, FL, USA		147		University of Florida
21		128.227.236.14	csev1-core-msfc-1-v41-1.ns.ufl.edu	Gainesville, FL, USA		156		University of Florida
22		128.227.254.74	-	Gainesville, FL, USA		146		University of Florida
23		**128.227.205.2**	cise.ufl.edu	Gainesville, FL, USA		131		University of Florida

You can find such "trace route" facilities by Googling, and then type in the IP-Addresses around the world

Route Across the US

- The route
  - Starts with my ISP moving packet through MAN
  - Next, the packet enters a regional WAN
  - Next, the packet crosses the backbone
  - Arriving at another regional WAN
  - Next arriving on campus in a LAN
  - Delivered to the destination computer

Regional Network

- The Watchtower regional network of Eastern MA
- Every state/region has one or a few

A Backbone Carrier -- NCSA

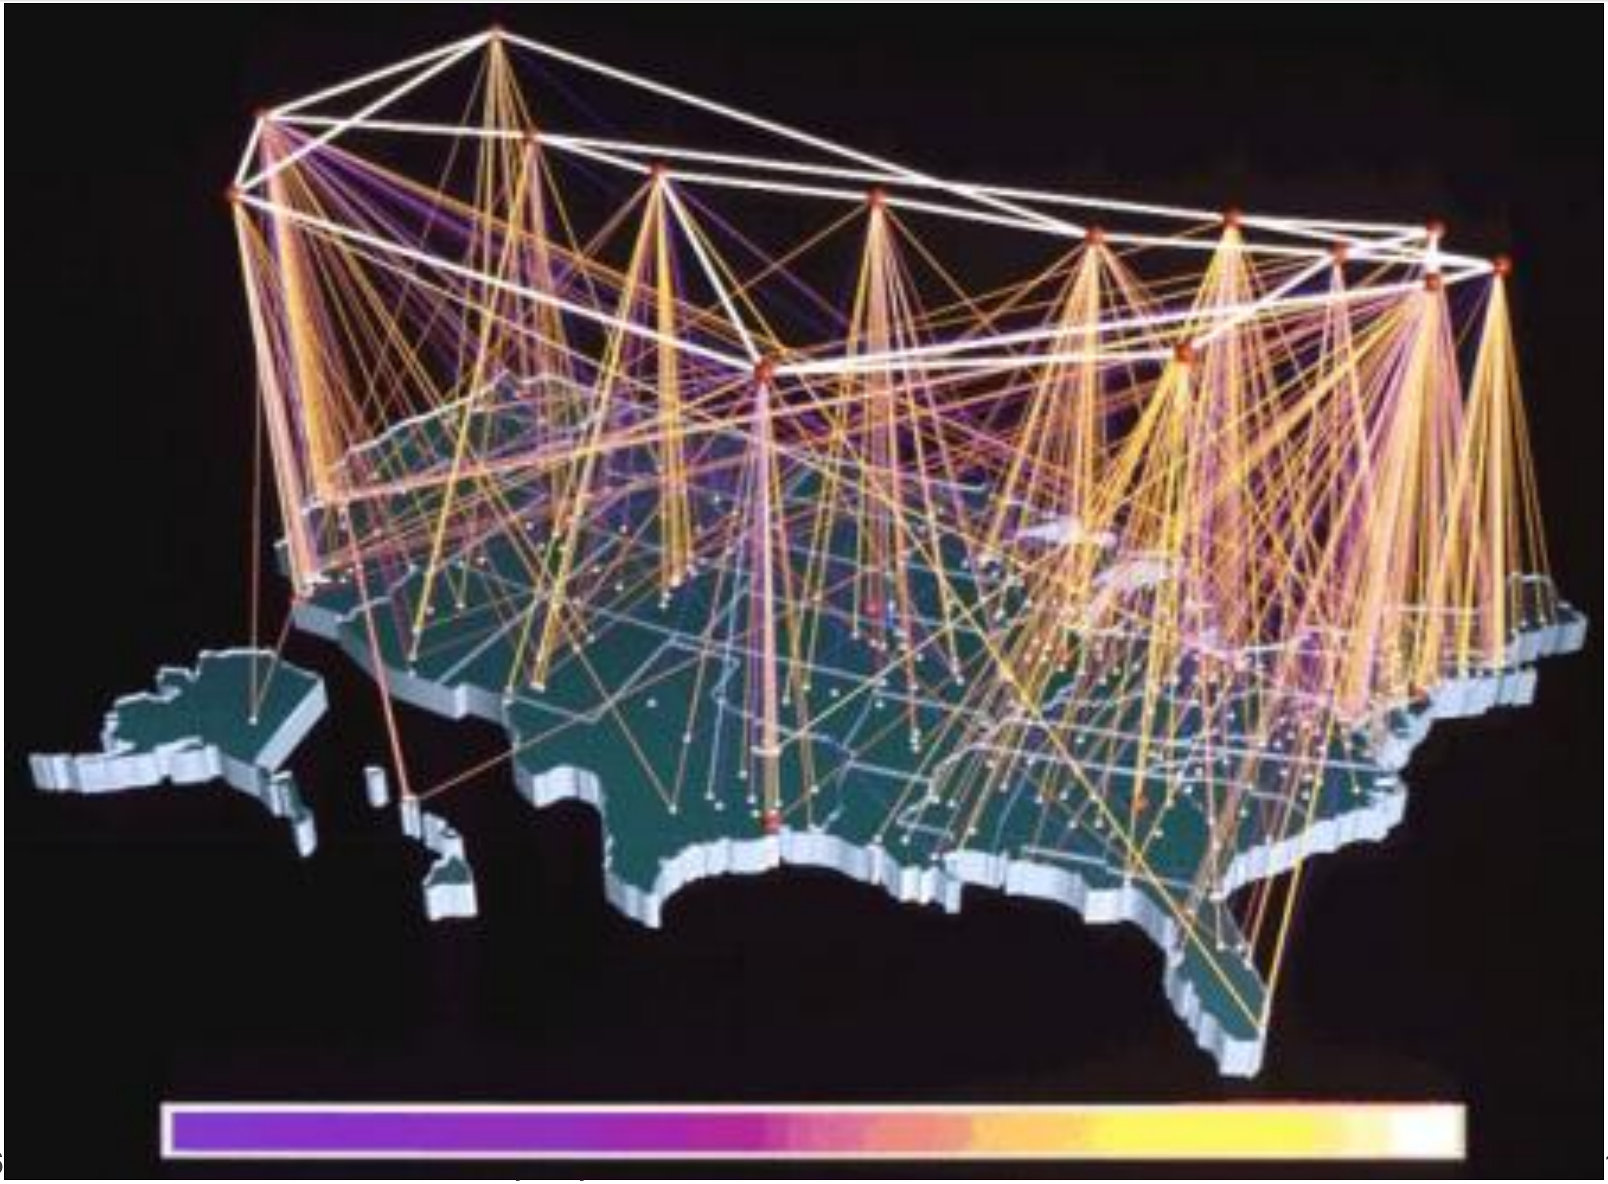

Wireless is a LAN technology

- As with "wired Ethernet," all computers in range can hear the radio signals of the others

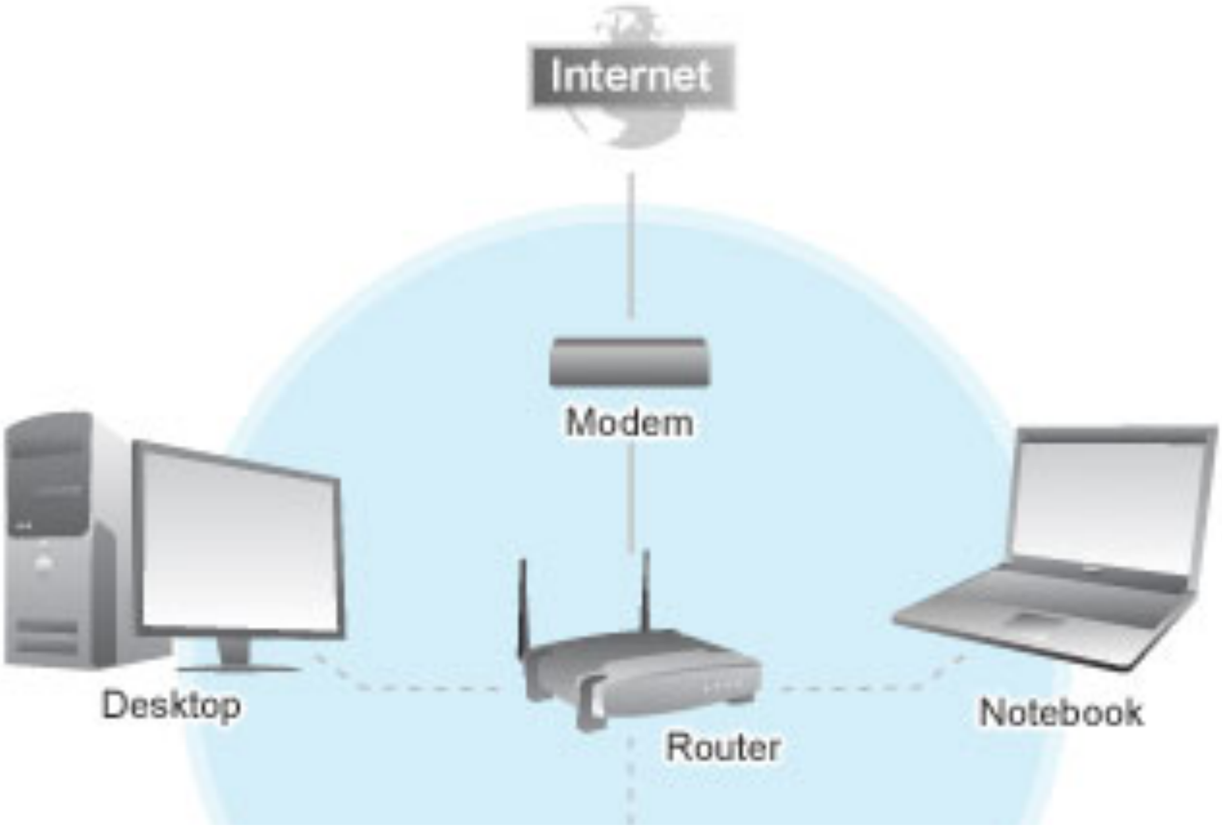

Naming Computers—Take 1

People name computers by a domain name

- a hierarchical scheme that groups like computers
  - .edu All educational computers, a TLD
  - .ucsc.edu All computers at UCSC
  - .ic.ucsc..edu All Instructional Computing computers
  - .soe.ucsc.edu School of Engr. Computers
  - riverdance.soe.ucsc.edu An SOE computer
  - unix.ic.ucsc.edu An IC computer.

Domains begin with a "dot" and get "larger" going right

Naming Computers—Take 2

Computers are named by IP address, four numbers in the range 0-255

unix.ic.ucsc.edu: **128.114**.104.230

riverdance.soe.ucsc.edu: **128.114**.48.104

- Remembering IP addresses would be brutal for humans, so we use domains
- Computers find the IP address for a domain name from the *Domain Name System*—an IP address-book computer

A computer needs to know IP address of DNS server!

Domains

.edu .com .mil .gov .org .net domains are "top level domains" for the US

- Recently, new TLD names added
- Each country has a top level domain name:
  - .ca (Canada)
  - .es (Spain)
  - .de (Germany)
  - .au (Australia)
  - .at (Austria)
  - .us (US)

Do you know sites like:
bit.ly
www.nba.tv
del.icio.us
… they exploit TLDs

Logical vs Physical

View the Internet in two ways:

1. Humans see a hierarchy of domains relating computers—logical network
2. Computers see groups of four number IP addresses—physical network

Both are ideal for the "user's" needs

- The Domain Name System (DNS) relates the logical network to the physical network by translating domains to IP addresses

True or false, Internet addresses and names are hierarchical?

A. Both are hierarchical
B. Neither are hierarchical
C. Just the names
D. Just the addresses

Internet vs. World Wide Web

- Many people mis-use the terms "Internet" and "World Wide Web"
- Let's get them right

Internet: all of the wires, fibers, switches, routers etc. connecting named computers

Web: That part of the Internet —web servers —that store info and serve Web pages and provide other services to client computers

One More Protocol: Client/Server

- The Web and much of the Internet services use the client server form of interaction

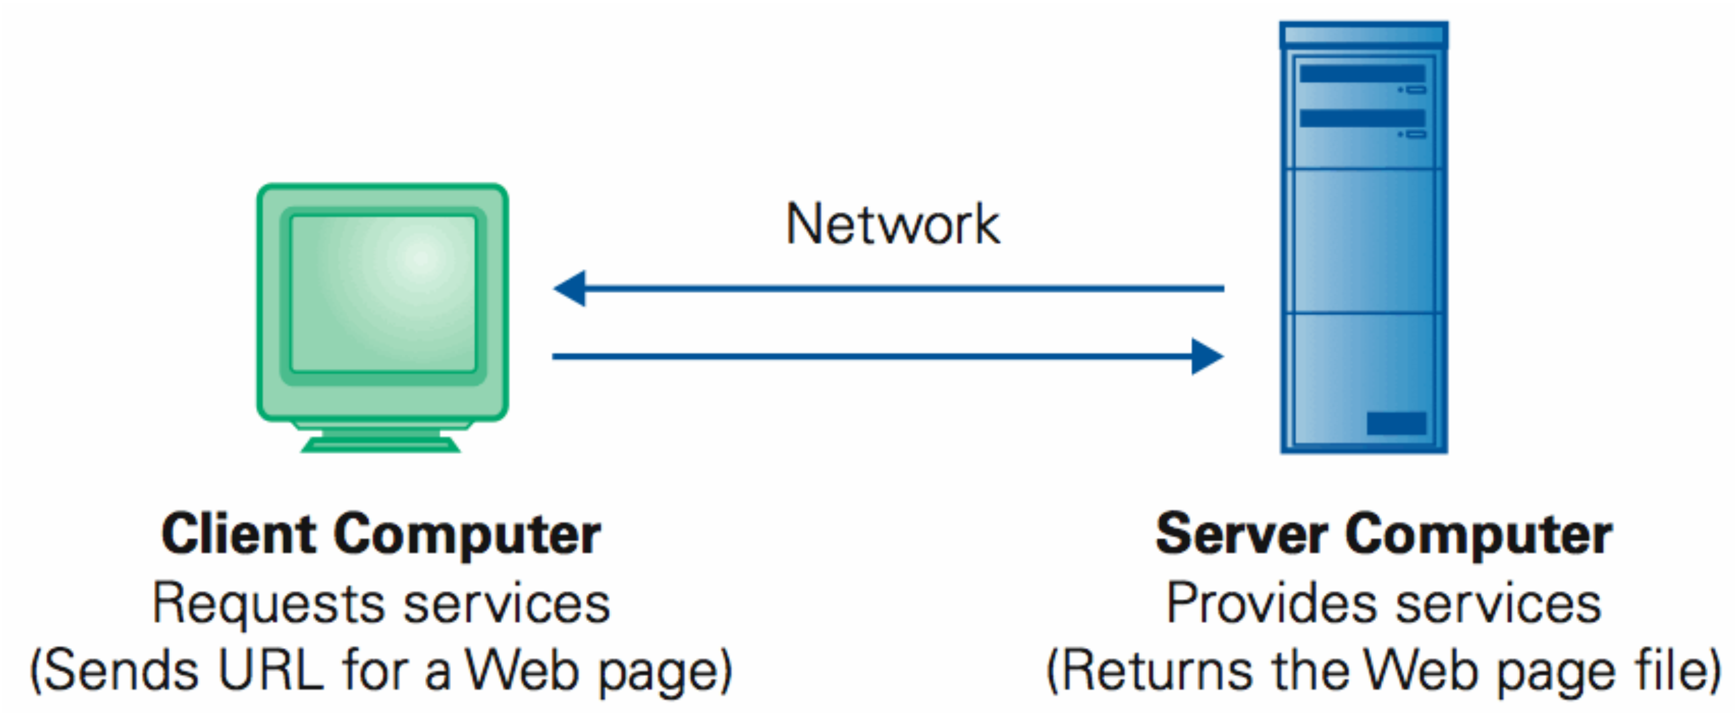

- It's a VERY BRIEF relationship

Client/Server Is Also Smart

- Clients and servers are not connected – they only exchange info … "no commitment issues"

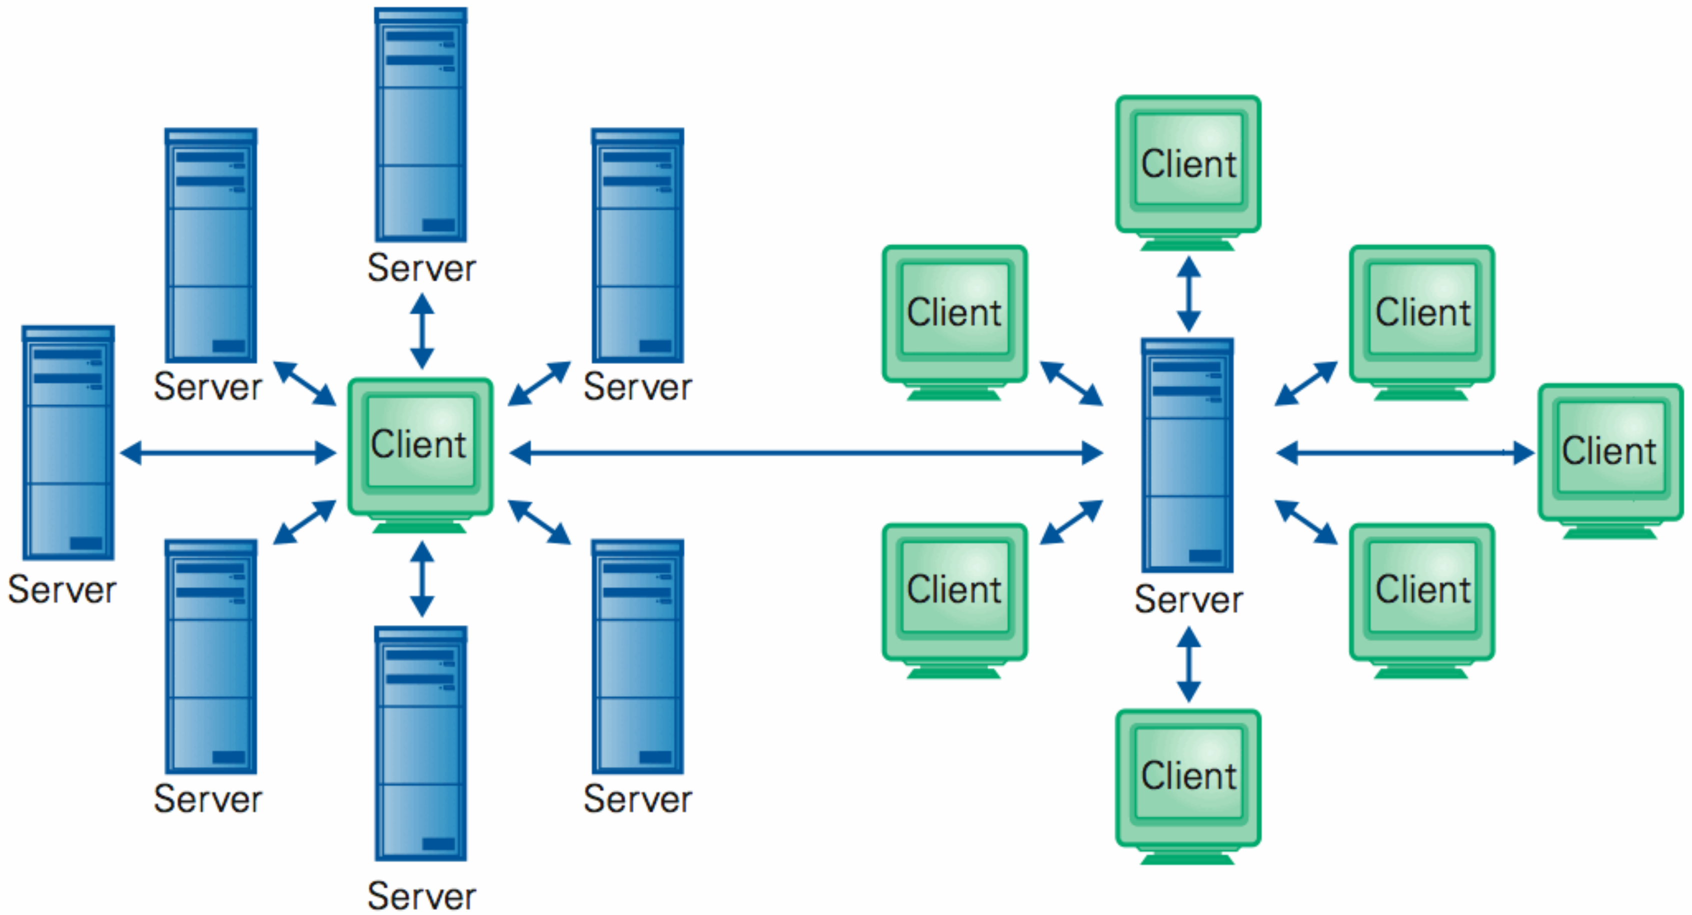

Internet ~~is Running~~ out of Addresses

HAS RUN

- The last top level IPv4 addresses were assigned in February 2011.
- IPv4 uses 32 bits and thus can specify 4.3 billion IP addresses.

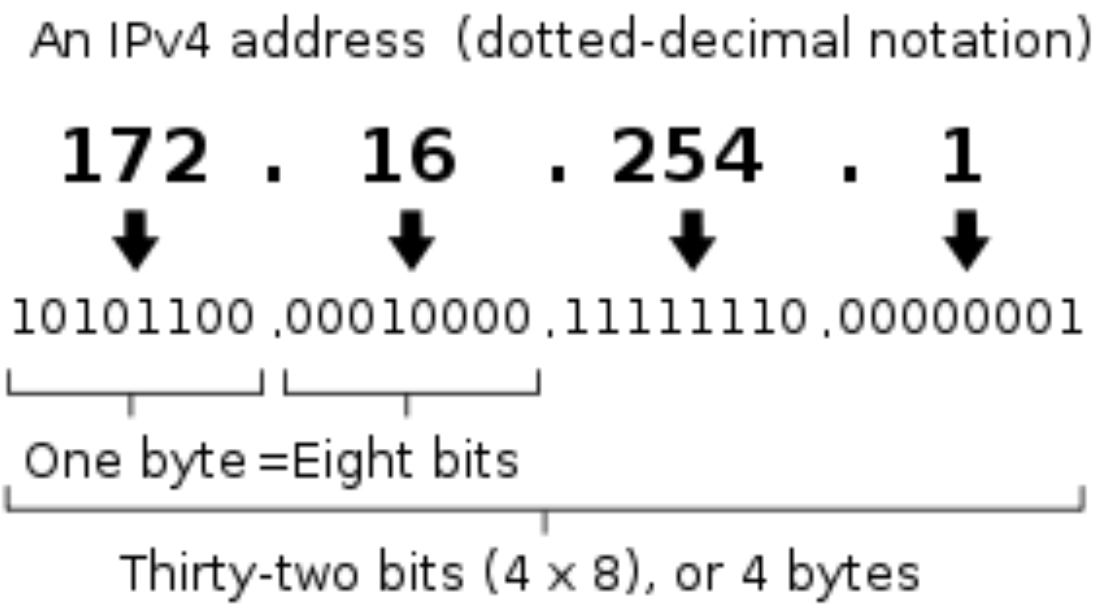

IPv6

The decision to put a 32-bit address space on there was the result of a year's battle among a bunch of engineers who couldn't make up their minds about **32, 128, or variable-length**. And after a year of fighting, I said—I'm now at ARPA, I'm running the program, I'm paying for this stuff, I'm using American tax dollars, and I wanted some progress because we didn't know if this was going to work. So I said: OK, it's 32-bits. That's enough for an experiment; **it's 4.3 billion terminations**. Even the Defense Department doesn't need 4.3 billion of everything and couldn't afford to buy 4.3 billion edge devices to do a test anyway**. So at the time I thought we were doing an experiment to prove the technology and that if it worked we'd have opportunity to do a production version of it. Well, it just escaped!** It got out and people started to use it, and then it became a commercial thing. So this [IPv6] is the production attempt at making the network scalable.
—Vint Cerf, *Google IPv6 Conference 2008*

IPv6

- 128 bit addresses
- Approximately 3.4×10^{38} addresses.
- That's 4.8×10^{28} addresses for each of the 7 billion people alive in 2011.

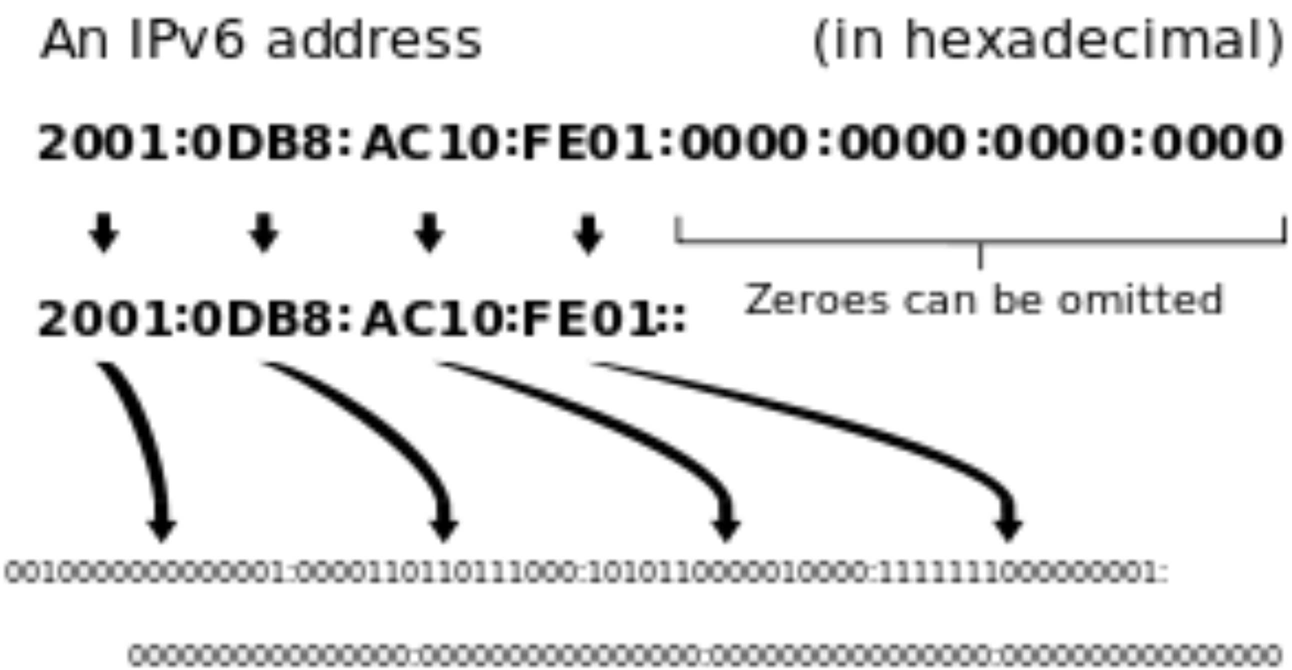

Summary

Networking changed the world

Internet: named computers using TCP/IP

WWW: servers providing Web pages

- Principles
  - Logical network of domain names
  - Physical network of IP addresses
  - Protocols rule: LAN, TCP/IP, http…
  - Domain Name System connects the two
  - Client/Server, fleeting relationship on WWW

Which item does not belong in this list?

A. Ethernet
B. ftp
C. http
D. html
E. X.25

Relating the “logical” with the “physical”

Domain Name System

Lawrence Snyder
University of Washington, Seattle

Recall 2 Ways To Name Computers

- **Logical**: Humans use domain names
  - riverdance.soe.ucsc.edu
- **Physical**: Computers use number-quads
  - 128.114.48.104
- This is different than the phone system:
  - The people use numbers: 1 800 555 1212
  - The equipment uses the same numbers
- A key property of computers: they can separate the logical form (preferred by people) from the physical form they must use

Today ...

- Today, we explain how the logical/physical separation is implemented for domain names
- But, this is also a chance to illustrate the structure of LARGE systems
  - Study the basic components
  - Study design ideas that make the system work well
  - This matters to you because you'll probably have "big ideas" about using computers

What's the Problem?

- The Internet is completely decentralized
  - No one is in charge – ICANN
  - A few companies get permission to give users or organizations IP-addresses – not much logic to it
  - When a person or organization gets an IP-address, it picks a domain name – few rules
- Once connected to I' net, users start using domain name … but when someone refers to it, how does their computer get its number?

Internet Corporation for Assigned Names and Numbers

Recall mail to "friend@cise.ufl.edu"

68.87.205.1	-	Mt Laurel, usa
68.85.240.101	be-70-ar01.burien.wa.seattle.comcast.net	Mt Laurel, usa
68.85.240.69	be-30-ar01.seattle.wa.seattle.comcast.net	Seattle, WA, USA
68.86.90.213	pos-0-5-0-0-cr01.seattle.wa.ibone.comcast.net	Seattle, WA, USA
68.86.85.206	pos-0-8-0-0-cr01.portland.or.ibone.comcast.net	Portland, OR, USA
68.86.85.197	pos-1-15-0-0-cr01.sacramento.ca.ibone.comcast	Sacramento, CA, USA
68.86.85.181	pos-0-9-0-0-cr01.sanjose.ca.ibone.comcast.net	San Jose, CA, USA
154.54.11.105	te3-3.mpd01.sjc04.atlas.cogentco.com	San Jose, CA, USA
154.54.0.177	te9-1.ccr02.sfo01.atlas.cogentco.com	San Francisco, CA, USA
154.54.3.137	te3-8.ccr01.lax01.atlas.cogentco.com	Los Angeles, CA, USA
154.54.0.226	te3-8.ccr01.iah01.atlas.cogentco.com	Houston, TX, USA
154.54.24.194	te3-2.ccr01.mia01.atlas.cogentco.com	Miami, FL, USA
154.54.1.186	te3-3.ccr01.mia03.atlas.cogentco.com	Miami, FL, USA
38.112.31.66	florida_lambdarail_llc.demarc.cogentco.com	Washington, DC, USA
198.32.155.10	tpa-flrcore-7609-1-te21-1.net.flrnet.org	Marina del Rey, usa
198.32.173.161	tlh-flrcore-7609-1-te41-1907.net.flrnet.org	Marina del Rey, usa
198.32.173.162	ctx36-ewan-msfc-1-v1907-1.ns.ufl.edu	Marina del Rey, usa
128.227.236.85	ctx36-nexus-msfc-1-v50-1.ns.ufl.edu	Gainesville, FL, USA
128.227.236.14	csev1-core-msfc-1-v41-1.ns.ufl.edu	Gainesville, FL, USA
128.227.254.74	-	Gainesville, FL, USA
128.227.205.2	cise.ufl.edu	Gainesville, FL, USA

- A packet sent to 128.227.205.2 finds its way

But, how do we get 128.227.205.2?

- When we send mail to a friend at the U of FL, we type friend@cise.ufl.edu and the computer that sends mail for us on campus needs to find out this fact:

cise.ufl.edu == 128.227.205.2

- We said it asks the Domain Name System, or DNS … so what happens

But Wait!

- How does it know the address of the DNS?
- You (or someone or something who set up your computer) told it when connecting it to the network … look in net control panel

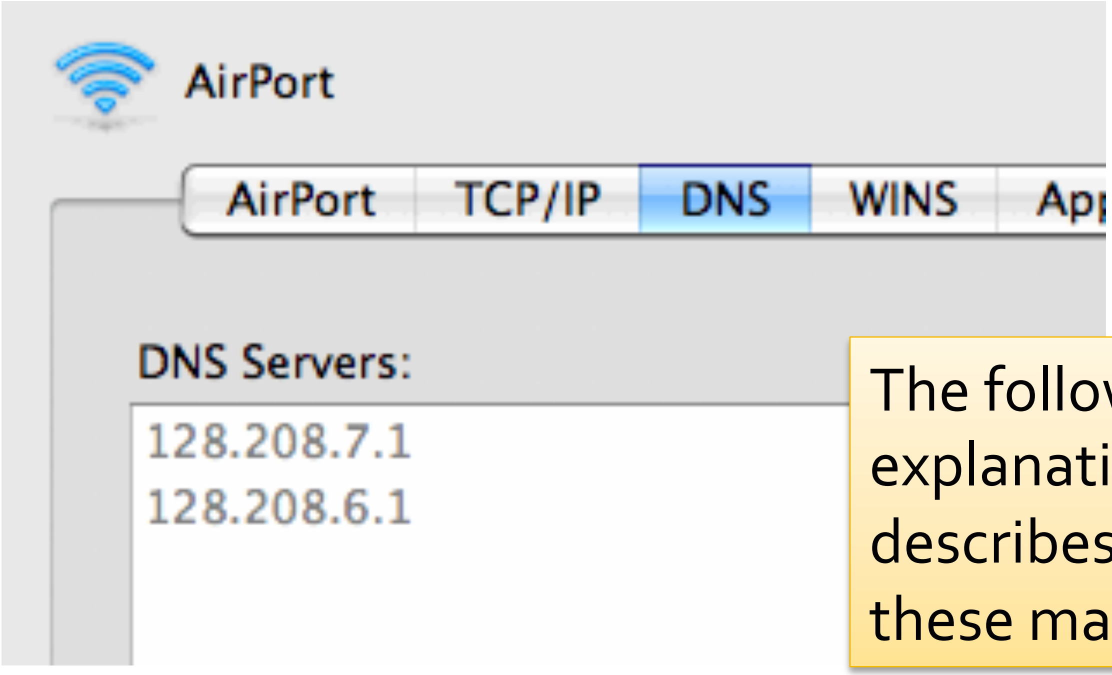

The following explanation describes what these machines do

First Step

- The DNS server answers the question "what number is cise.ufl.edu?" by this method
- First Step: Look it up in the "address book"
  - The DNS server does that
  - It keeps its own address book, a list of all of the domain names like cise.ufl.edu that it has been asked about and found
  - We say it *caches* the addresses it's found
    - *caching – keeping a copy around in case need it again*
  - It checks the cache first

If It Has Never Been Asked …

- The address will not be in the cache if this is the first request
- Second Step: The DNS server begins a process of finding the address on behalf of your computer …

  That process uses 2 Facts of I' net

The DNS Design: Fact 1

- Every domain has an authoritative name server, which I' ll call autho

eenie 123.4.5.3
meenie 123.4.5.2
minie 123.4.5.1
mo 123.4.5.0

- Two Cases: Autho knows the number of every computer in its domain

The DNS Design: Fact 1 (Continued)

- OR Autho knows the number of every autho computer in its domain

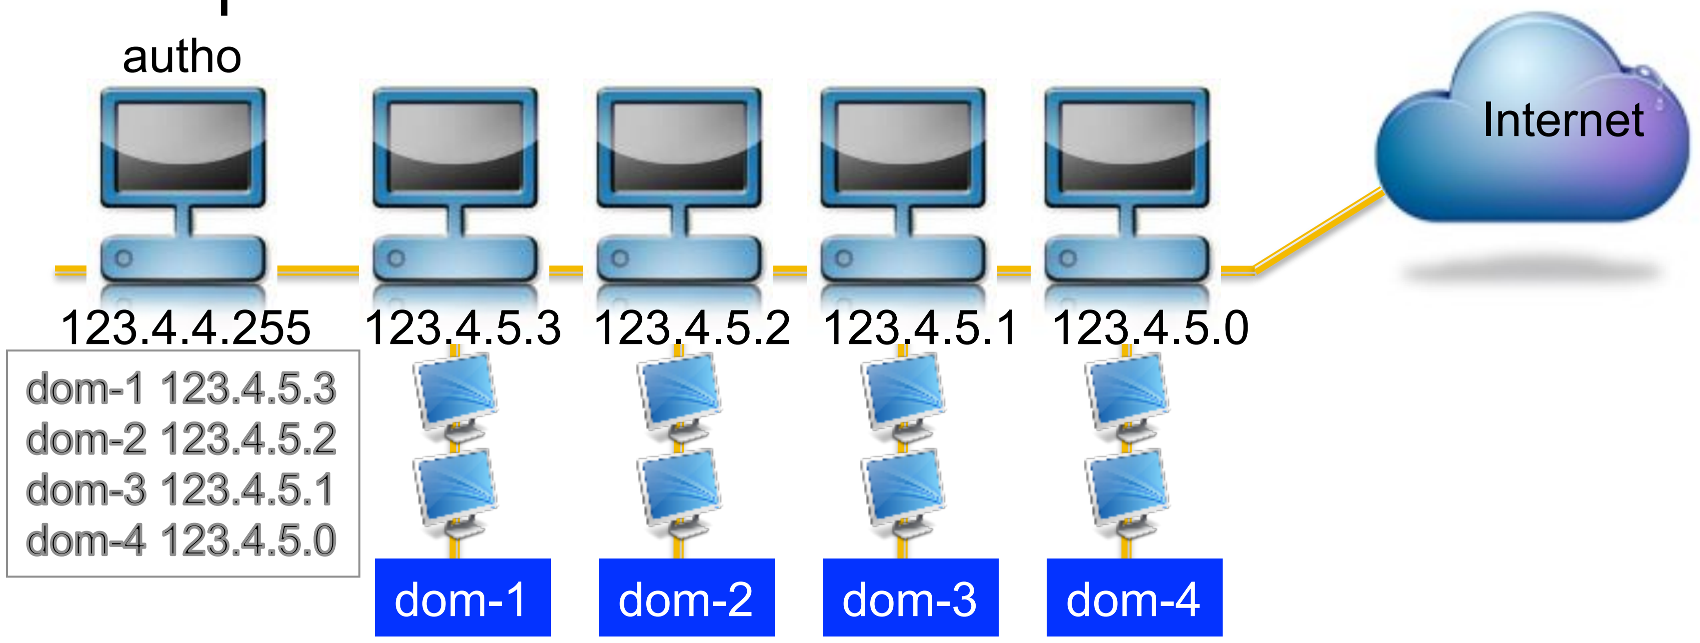

The DNS Design: Fact 2

- There are 13 Internet “root name servers” scattered around the world … all the same

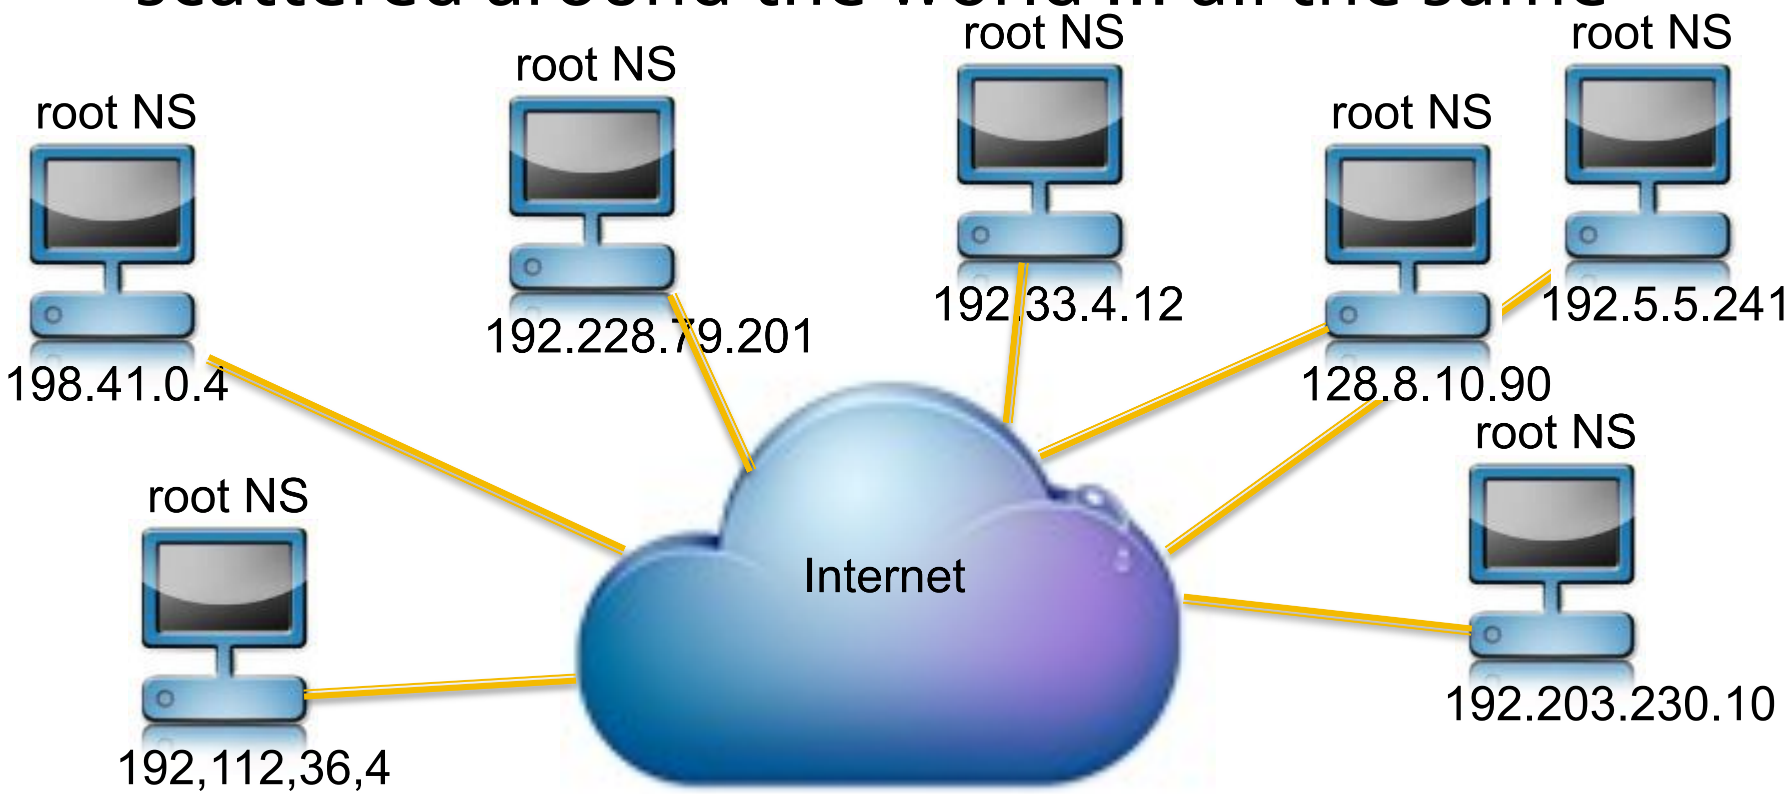

- All DNS servers have their numbers

So, Here's How It Goes …

- Your computer's DNS server never heard of cise.ufl.edu.root … so it pulls the domain name apart:
  - cise, a computer in the .ufl domain
  - ufl, a domain in the .edu domain
  - edu, a domain in the .root domain
- So, the DNS begins at the end and starts asking for the numbers of the autho computers … who's the autho for the .root domain?

Your DNS Asks the .root NS

- Please give me the number of .edu autho
  - Getting that it asks it, …
- Please give me the number of .ufl autho
  - Getting that it asks it, …
- Please give me the number of the cise machine
  - Getting 128.227.205.2, it addresses your email and sends it on
- Simplification: it might have cached .edu autho and .ufl autho, which saves those requests

Logical Names Form Hierarchy

- As a hierarchy, it can be shown as a tree:

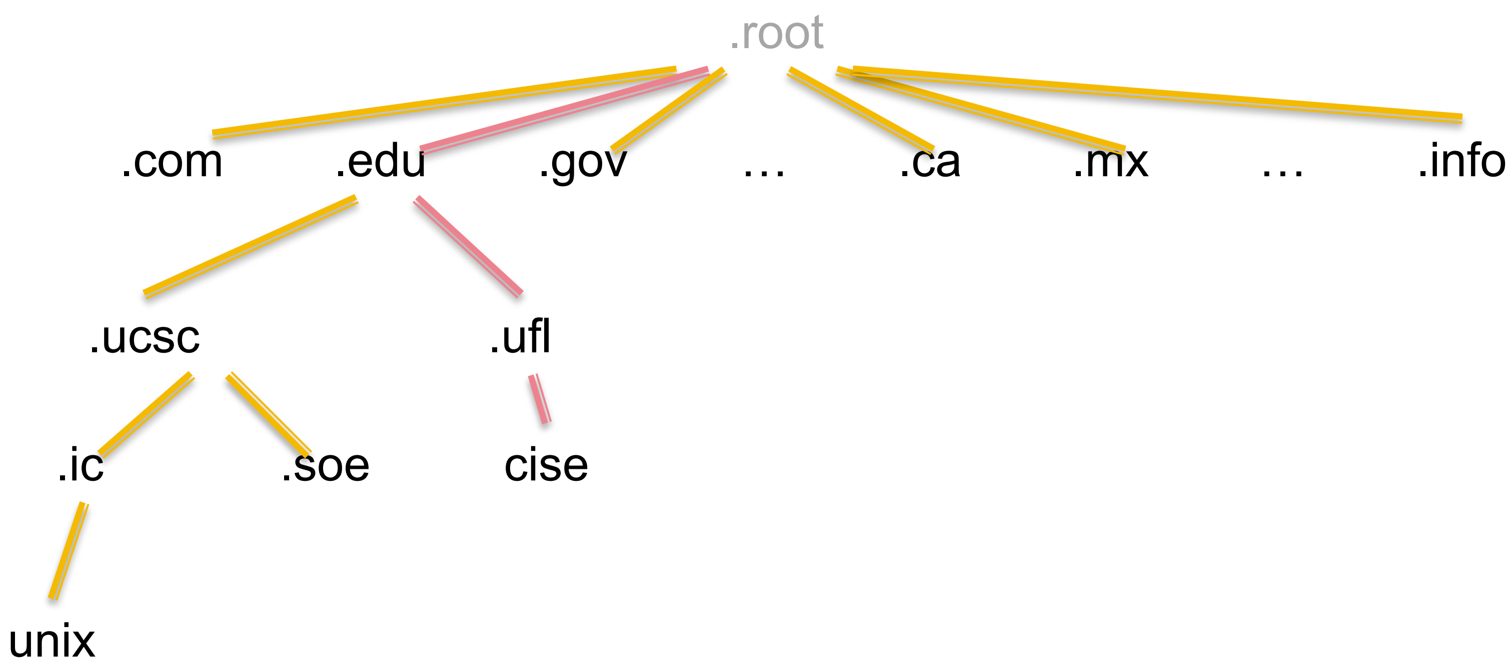

- The DNS is simply "walking" down the tree asking each autho for the number of next item

Exercise:

- I was in Minneapolis in October working at a hotel and went to log into my computer at UCSC
  - riverdance.soe.ucsc.edu
- How did the hotel's ISP find 128.114.48.104?

Finding riverdance.cse.ucsc.edu

The maximum number of autho's that the Baltimore hotel's DNS might have needed to consult is

A. 1
B. 2
C. 3
D. 4
E. 5

Finding riverdance.cse.ucsc.edu

The minimum number of autho's that the Baltimore hotel's DNS might have needed to consult is

A. 0
B. 1
C. 2
D. 3
E. 4

Think About This Scheme: Huge

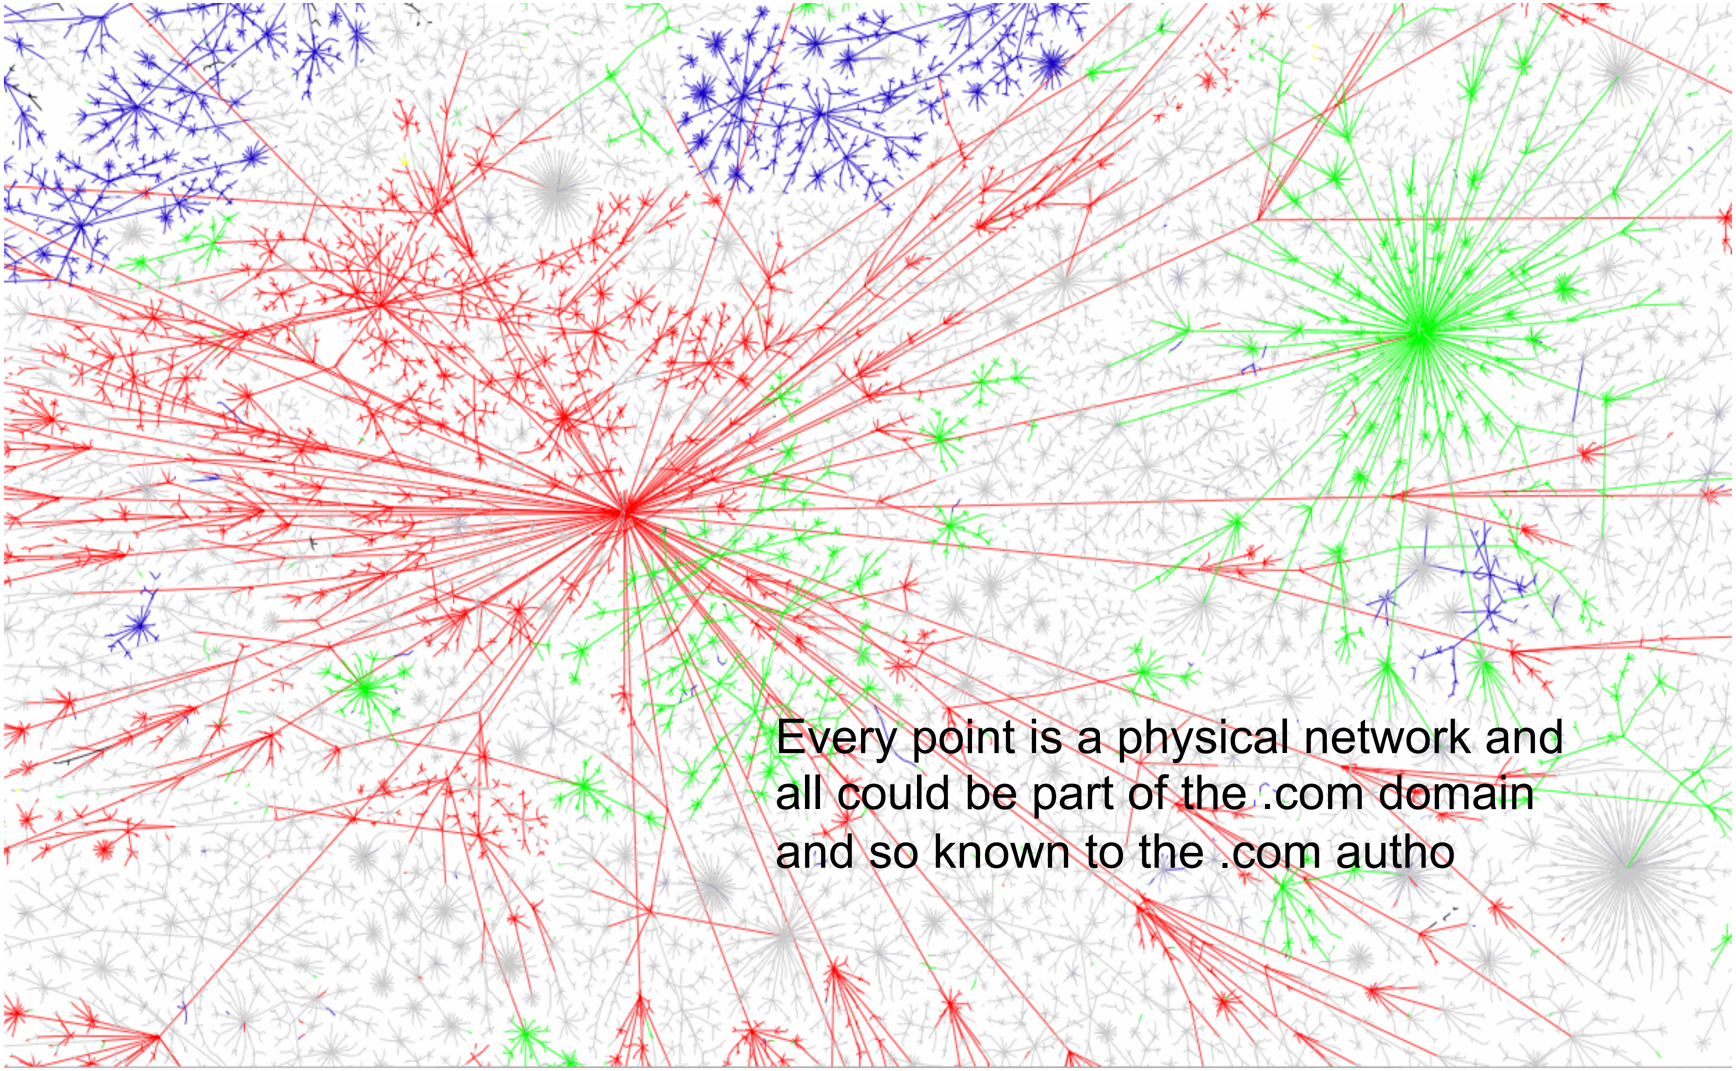

Suppose A Domain Adds Computer

- When a domain, say .ufl, adds a new computer it gets a name and an IP-address
- They add its name and number to the list in ufl autho's memory and its up and running, "known to the world"
- This is a completely decentralized solution – no one needs to be in charge except to make sure that the domain autho is up & correct

Properties …

- Fault tolerant: when a hurricane takes out Miami's power, only the domains without power are affected …
- Robust: when a fire burns down the building of a .root name server, 12 others can carry the load
- Enormous capacity: most lookups are independent and do not collide (b/c higher level domain authos are cached), but more capacity is possible by replicating authos

Compare DNS Structure To ..

- Master List Solution …
  - Suppose the design was for the root NS computers to have a master list of all

    domain_name: IP-address

    pairs connected to the Internet
- How would it be different, better or worse?