Networks – The Big Picture

• Types of "networks" and their characteristics • The Internet – Components, Structure • The Internet – Routing, Switching, Naming, Applications & Services • Historic Perspective – Design goals, E2E Principle • What are major evolution points – middleboxes (e.g., firewalls, NATs), addressing (CIDR), services (Akamai versus normal) • How did the application space evolve? Client-server, peer-to-peer, service-based (cloud computing), push computing? • What are major issues – failures, stability, performance (delay tolerance), mobility, attacks • "Clean-Slate" Designs.. What are they?? Why are they needed?

Networks Everywhere!

- Networks occur in many situations
  - ***Physical networks***: Internet, transportation networks, etc.
  - ***Biological networks***: Gene regulatory networks, protein activation networks, etc.
  - ***Social networks***: Online social networks, co-authorship/citation networks, blogosphere networks, terrorist networks, etc.

Network Science

- Network Science
  - Motivated by the omnipresence of networks in complex problems
  - Emerging cross-discipline
- How is Network Science related to Internet?
  - Internet is also a complex network – although it is an engineered network
  - Network Science can help in more resilient network designs
  - For example, what designs can provide better fault (attack) tolerance?
  - The Internet is becoming a substrate (backbone) for large number of "overlay" networks
    - Blogosphere networks
    - Online social networks

Network Science...

- Major thrust of Network Science is to characterize networks (many ways)
  - Degree distributions
  - Node clustering
  - Hierarchical structure
- ***Degree distribution***: defines how the degrees (e.g., out degree) of nodes are distributed

Network Science...

- Node degree distributions:
  - Power law distributions for many naturally occurring and Internet graphs

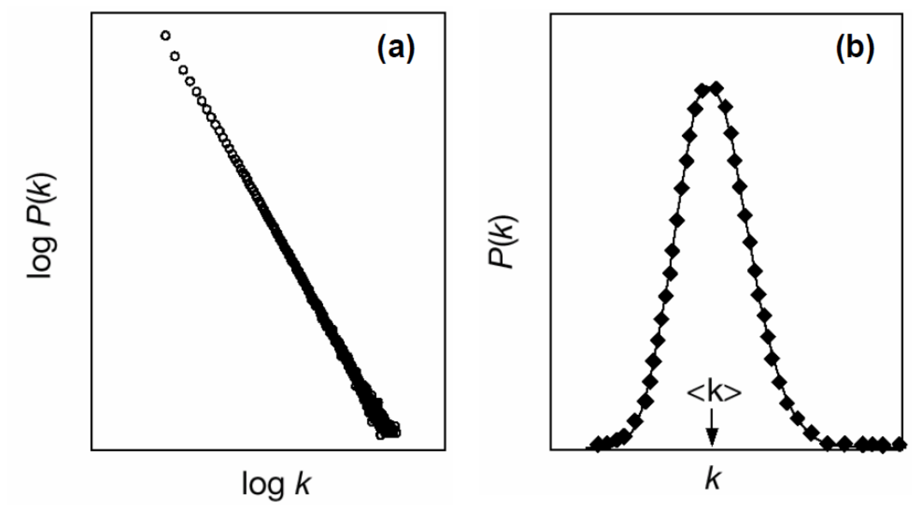

  - (a) sho [illegible] egree nodes and many low degree nodes (e.g., Internet)
  - (b) “peeked distribution” (e.g., distribution of human heights) – average representative value for the distribution is available

Network Science...

- Why is it important to know the network model of the Internet?
  - Useful in simulation studies where we need to recreate the network structures
  - Understanding the failure properties of the Internet. For example, power-law graphs can be resilient to random failures and vulnerable to targeted attacks.

Network Science...

- Different network types

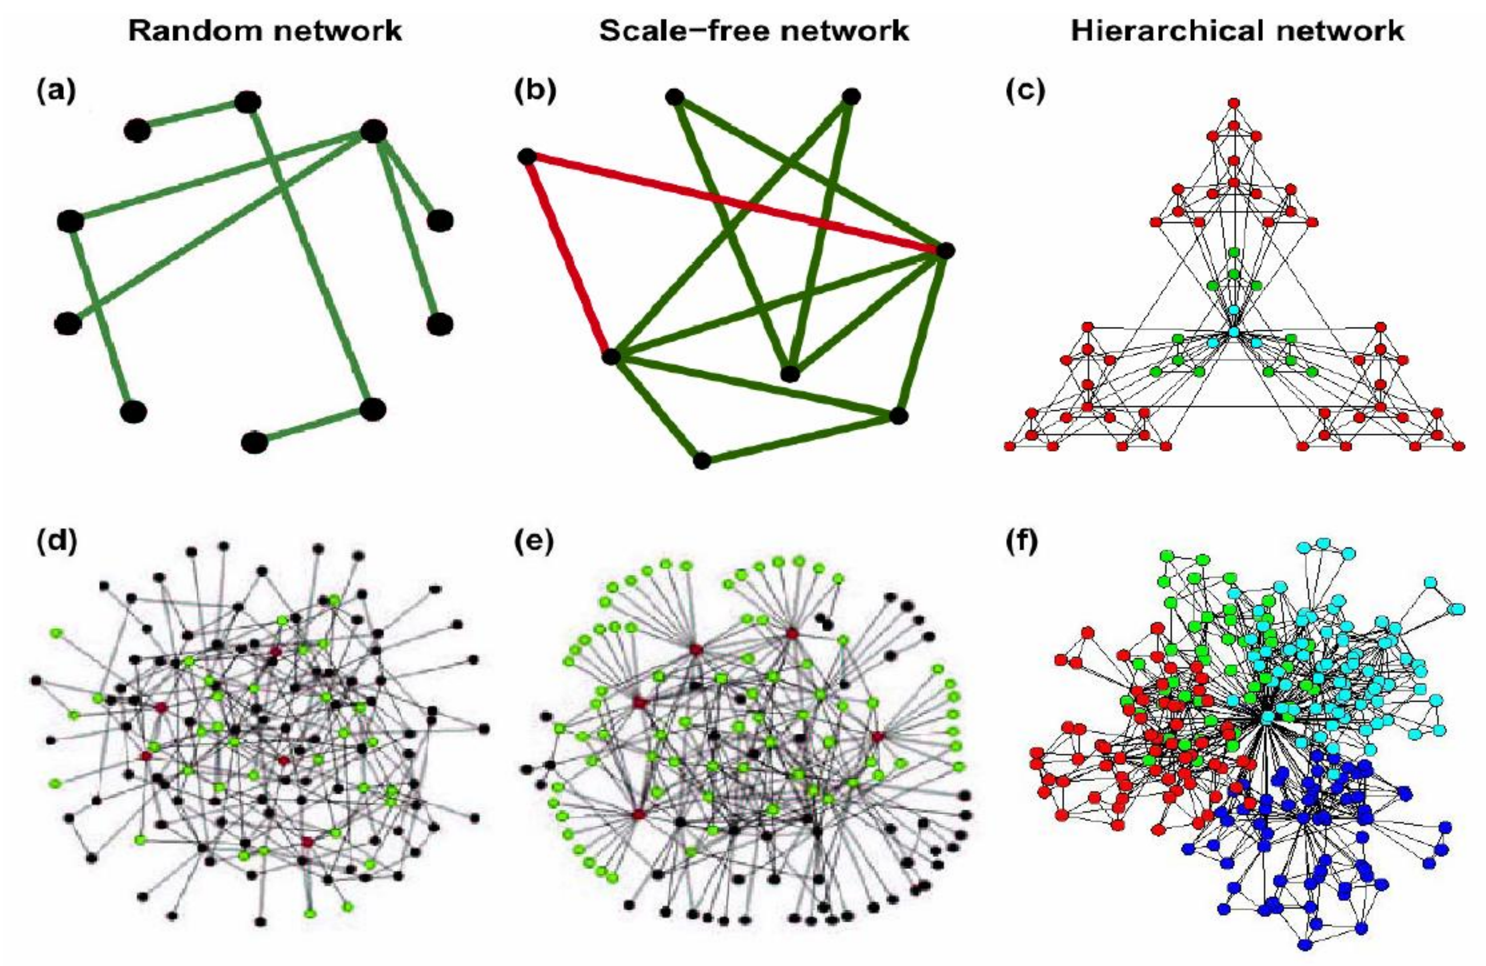

Lets Turn to Internet!

- A simple “nuts & bolts” view:
  - Routers and hosts as nodes and physical connections as edges (for wired networks)
  - Routers at “interior” nodes
  - Hosts at “exterior” or leaf nodes

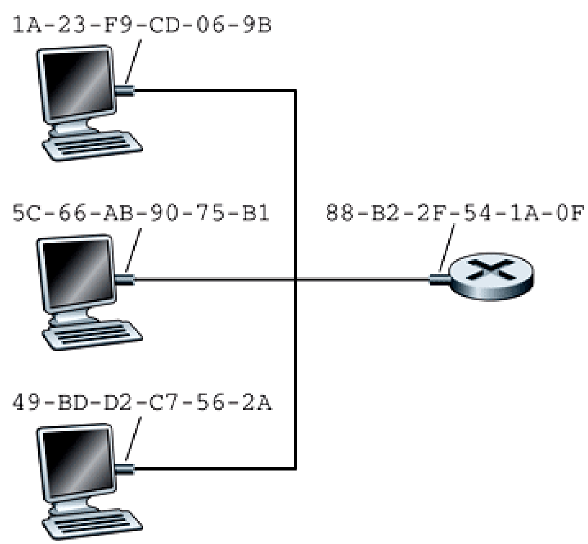

A simple LAN

Internet Structure

- Internet is a “network” of networks
  - Multiple independently owned and operated networks are connected to provide end-to-end connectivity
- Different types of networks:
  - Company or campus networks
  - Internet service providers (regional or global)
  - Home networks
  - Mobile networks (increasing in scope & size)

Internet Structure…

Wide area wireless, Vehicular networks

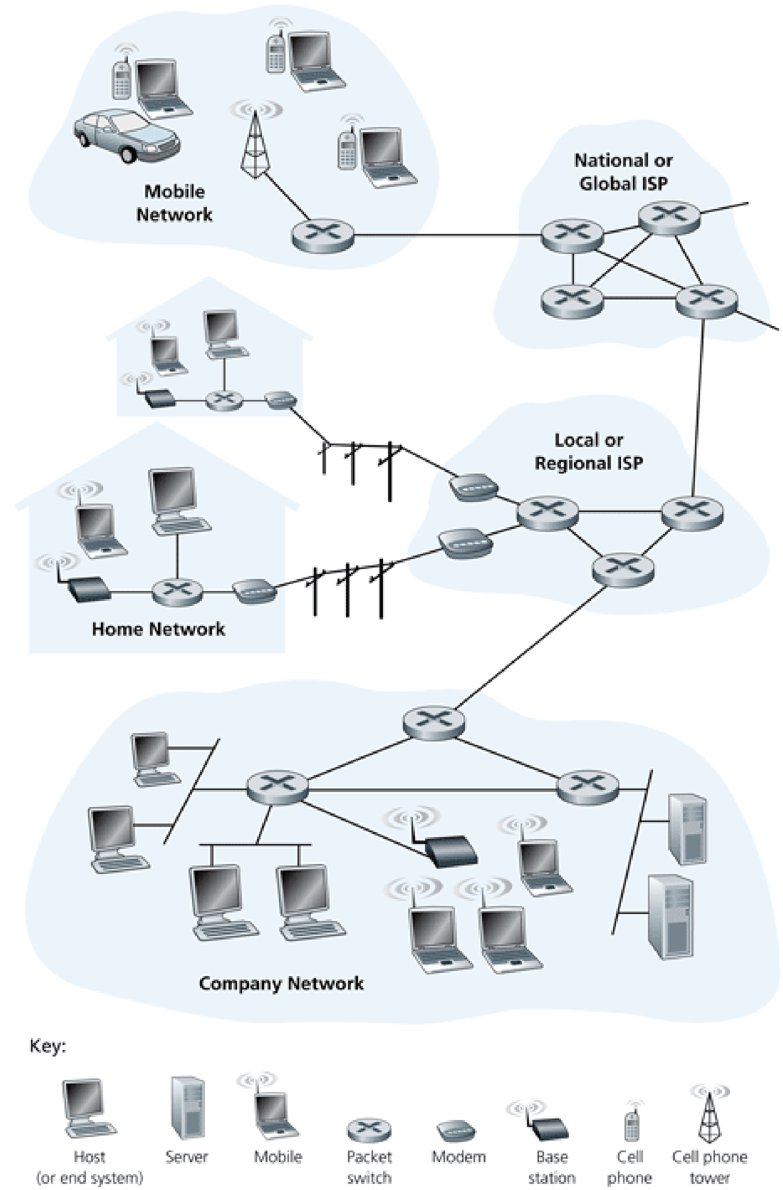

Internet Structure…

- Independently managed networks called Autonomous System (AS)

AS	Name	ISP		with customer & peer		POPs
		Routers	Links	Routers	Links	
1221	Telstra (Australia)	355	700	2,796	3,000	61
1239	Sprintlink (US)	547	1,600	8,355	9,500	43
1755	Ebone (Europe)	163	300	596	500	25
2914	Verio (US)	1,018	2,300	7,336	6,800	121
3257	Tiscali (Europe)	276	400	865	700	50
3356	Level3 (US)	624	5,300	3,446	6,700	52
3967	Exodus (US)	338	800	900	1,100	23
4755	VSNL (India)	11	12	121	69	10
6461	Abovenet (US)	367	1,000	2,259	1,400	21
7018	AT&T (US)	733	2,300	10,214	12,500	108

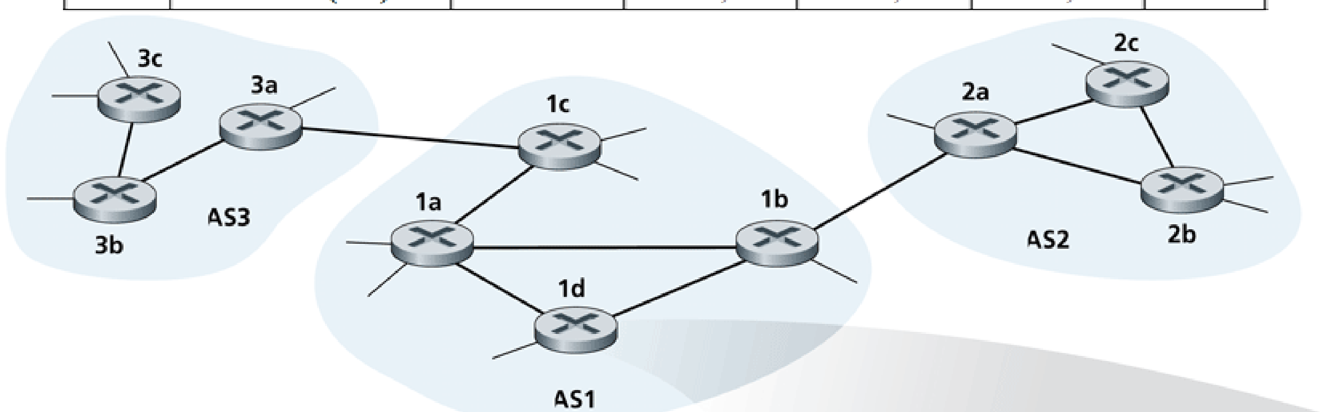

Internet Structure…

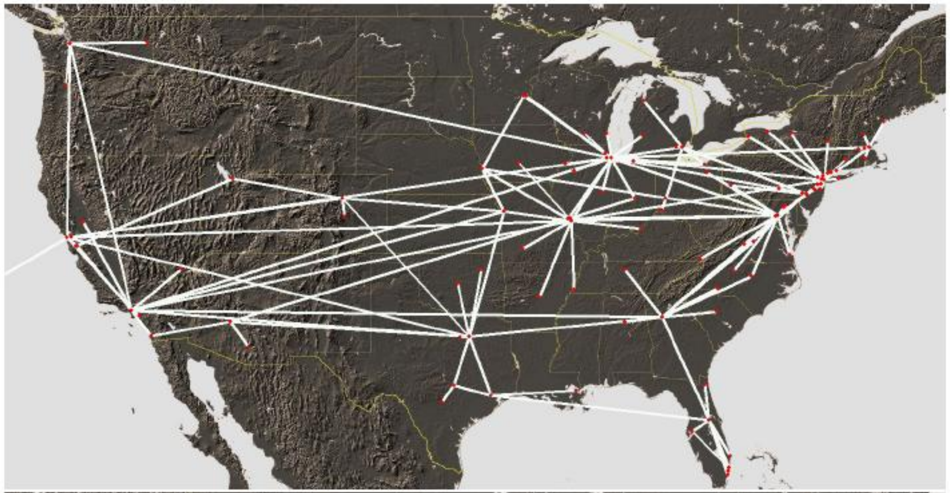

Backbone topology of AT&T

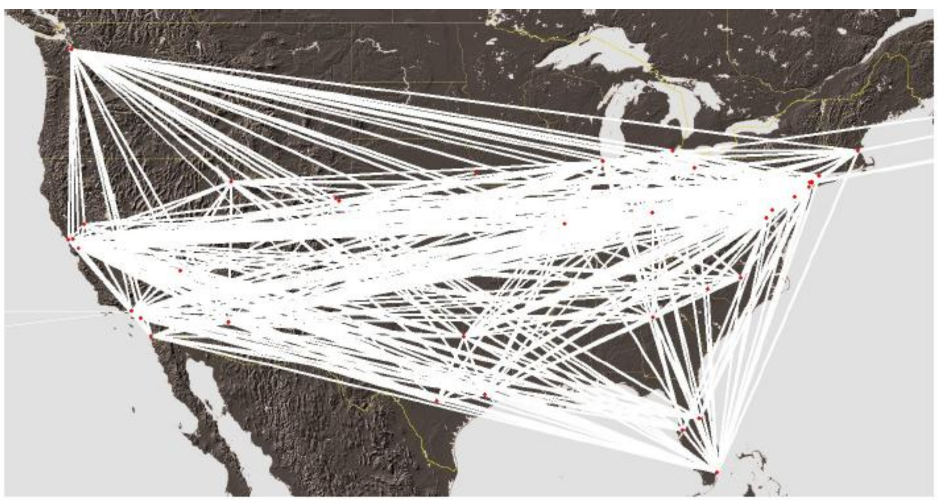

Backbone topology of Level 3
High connectedness can be due to virtual circuits

Topologies inferred by sending probes from selected points on the network

Switching and Routing

- How can packets “hop” through the network from *source* to *destination??*
  - ***Circuit switching*** – Establish a circuit before the packets are sent out (like a phone call)
  - ***Packet switching*** – Let the packets undergo a hop-by-hop distributed process to reach the destination
- *Packet switching* – One of the cornerstones of the Internet architecture – a deviation from preceding networks
  - phone networks
  - high-performance computer networks

Packet versus Circuit Switching

- With circuit switching a channel is occupied by a traffic flow even when the channel is not actively used
- With packet switching, the channel is multiplexed between flows
- Can get better utilization

Routing versus Forwarding

- Forwarding
  - Happens at each device – a process that determines how an incoming packet should be treated (specifically where to send)
  - Uses a local database (route table)
- Routing
  - A global process that aggregates the topological information to create a local database (route table) to be used by forwarding

Routing Issues on the Internet

- Simple approach to routing would be to use Dijkstra's algorithm to compute shortest paths
- Problems:
  - Scale of the Internet
  - Constant change of state – convergence speed, stability, stale routes
  - Multiple operators – do not share topological information (business sensitive information)
  - Traffic cross different networks at peering points where traffic is carefully ***metered***

Historic Perspective

- Internet started as a DARPA project
- Initial design goals:
  - Top-level: “develop an effective technique for multiplexed utilization of existing interconnected networks”
  - Second-level:
    - survive loss of networks and gateways
    - support multiple types of communication services
    - accommodate variety of networks
    - distributed management of network resources
    - architecture must of cost-effective

Historic Perspective

- What is the Internet architecture??
- Most argue there is none. See RFC 1958.
- At a high level, Internet community believe in the following:
  - Goal is connectivity
  - Tool is Internet protocol
  - Intelligence is end-to-end

End-to-end Principle

- “If a function in question can completely and correctly be implemented only at the end-points of the communication system; providing this function as a feature of the communication system itself is not possible”
- Example: Reliable data service between two nodes

End-to-end Principle

- Calls for a simple network
- Dumb network/Intelligent end system
  - Internet
- Intelligent network/dumb end system
  - Phone network
- The Internet is moving more towards the Intelligent network/dumb end system model!
- Why??

Evolution of the Simple Internet

- Simple Internet
  - End systems have globally unique addresses
  - Network does simple packet transportation from source to destination
- Many issues have motivated the evolution:
  - Shortage of addresses
  - Requirement to protect networks
  - Requirements to accelerate performance
  - Different stakeholders in the networking scene wanting to make profit!

Addressing Issues

- IPv4 used a class-full addressing scheme (originally)

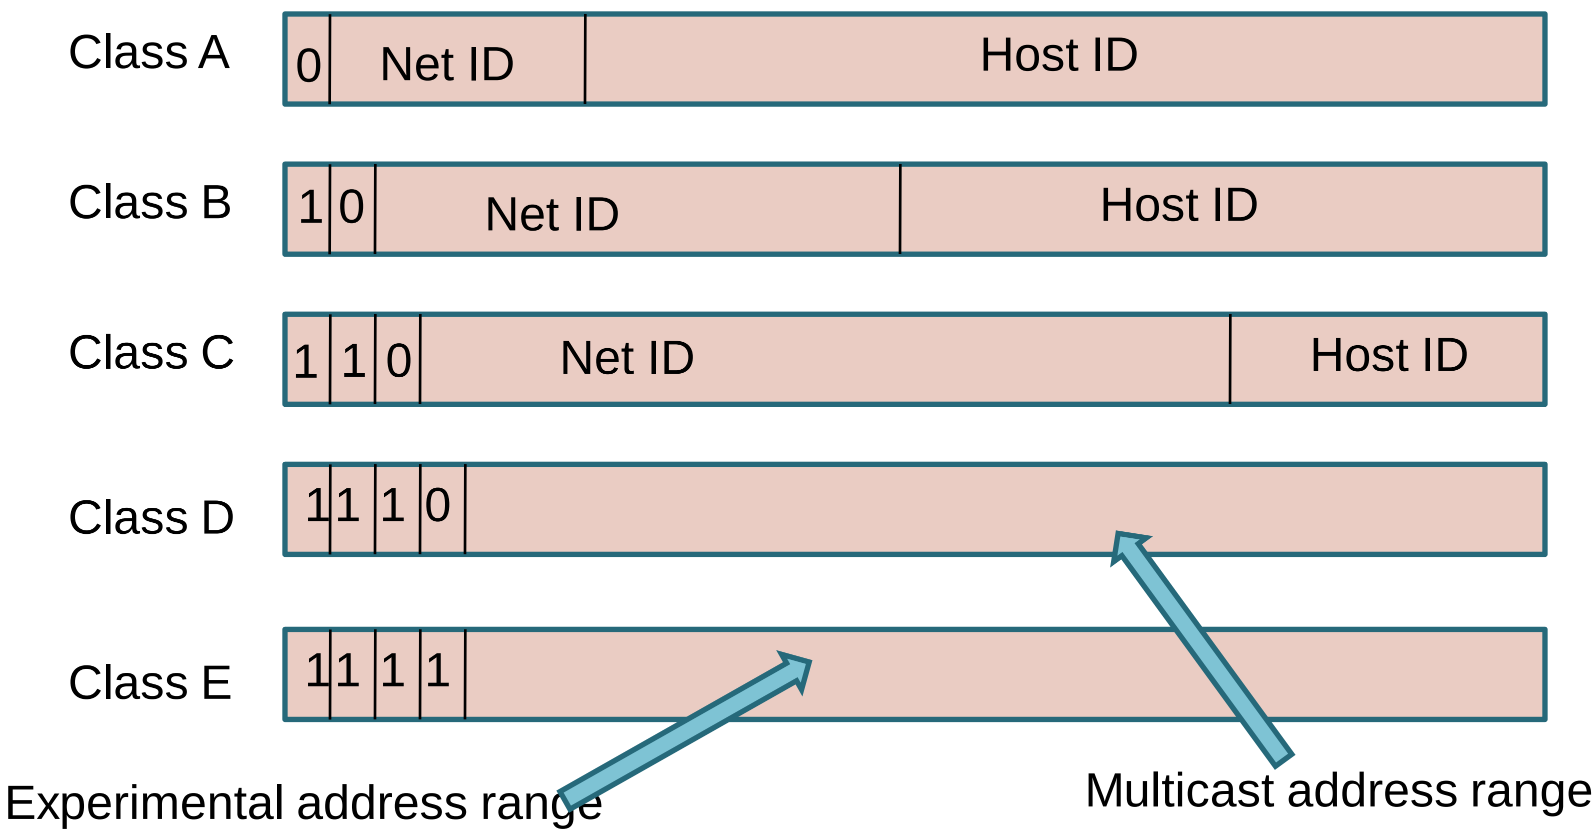

Addressing Issues...

- There are about 3,700 million usable addresses (managed in a ***class-less*** way)
- Table below shows regional address block (2^{24} addresses) assignments

Delegated to	Blocks	+/- 2007	Addresses (millions)
AfriNIC	2	+1	33.55
APNIC	26	+7	436.21
ARIN	27		452.98
LACNIC	6	+2	100.66
RIPE NCC	26	+4	436.21
Various	49	-1	822.08
End-user	42	-1	704.64
Available	43	-12	721.42

About 69% utilization of the IPv4 address space

Addressing Issues...

The 2568.14 million addresses currently in use aren't very evenly distributed over the countries in the world. The current top 15 is:

			2008-01-01	2007-01-01	change	Country
1	-	US	1408.15 M	1366.53 M	+3%	United States
2	-	JP	141.47 M	151.27 M	-6%	Japan
3	(4)	CN	135.31 M	98.02 M	+38%	China
4	(3)	EU	120.35 M	115.83 M	+4%	Multi-country in Europe
5	-	GB	83.50 M	93.91 M	-11%	United Kingdom
6	-	CA	73.20 M	71.32 M	+3%	Canada
7	-	DE	72.46 M	61.59 M	+18%	Germany
8	-	FR	67.79 M	58.23 M	+16%	France
9	-	KR	58.86 M	51.13 M	+15%	Korea
10	-	AU	33.43 M	30.64 M	+9%	Australia
11	(12)	IT	24.04 M	19.14 M	+26%	Italy
12	(11)	BR	23.46 M	19.27 M	+22%	Brazil
13	(16)	MX	21.50 M	16.26 M	+32%	Mexico
14	(13)	ES	20.42 M	18.69 M	+9%	Spain
15	-	NL	19.89 M	18.08 M	+10%	Netherlands

Middleboxes

- Middleboxes are transparently interposed between two communicating end-points
- Breaks few assumptions:
  - Globally unique and immutable addresses
  - End-points are aware of all changes
- Why middleboxes??

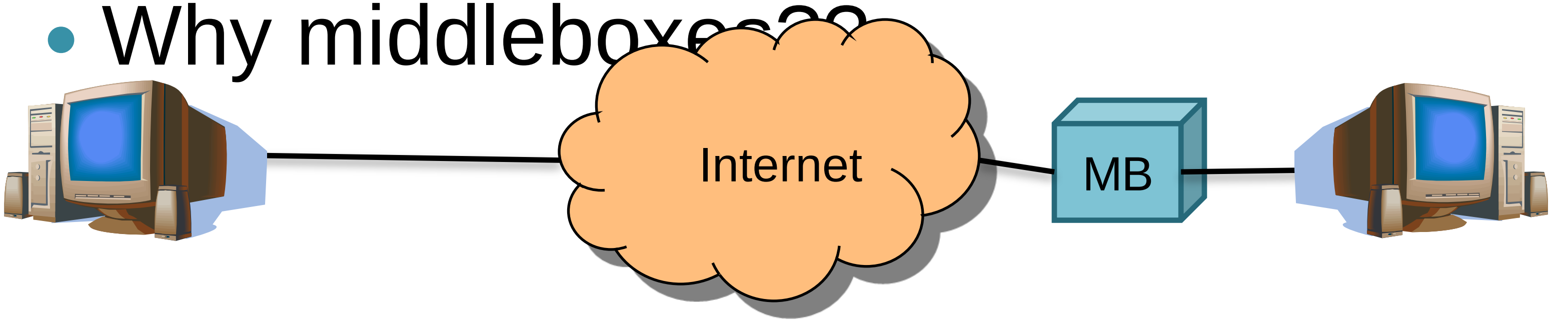

Middleboxes

- Various examples of middleboxes
  - Firewalls
  - Application accelerators
  - Transparent Proxy caches
  - Network address translators
  - Intrusion detection systems
- Middleboxes make the network complex and new application deployment harder

Wide-Area Services

- Client-side approaches for improving performance (proxy caching)
  - Pull-based technique
  - No way for the server to throttle the process

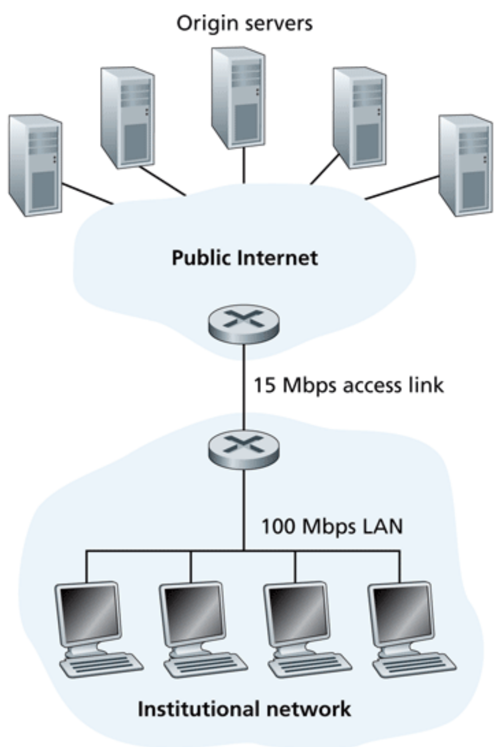

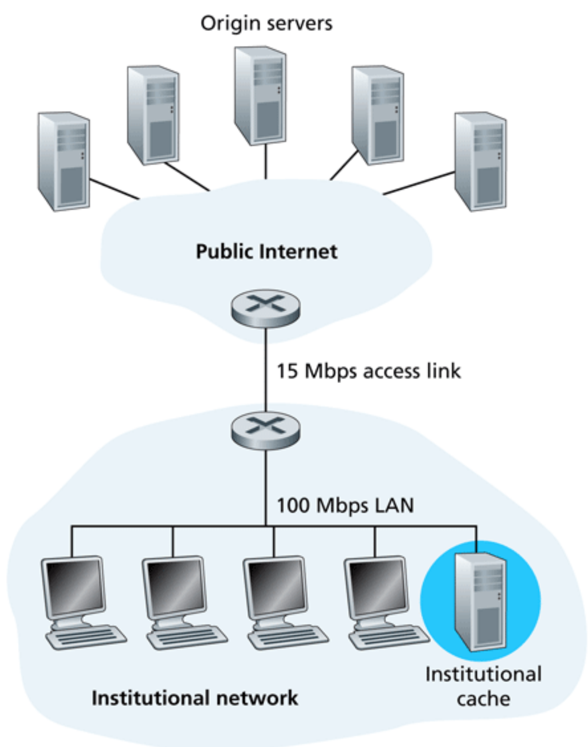

Wide-Area Services...

- Businesses want the ability to “place” their services above their competitors
- Push-based (content distribution) to ensure performance
- Uses variety of techniques:
  - Content/URL rewriting
  - Proxy placement
  - DNS based location aware load balancing

Wide-Area Services…

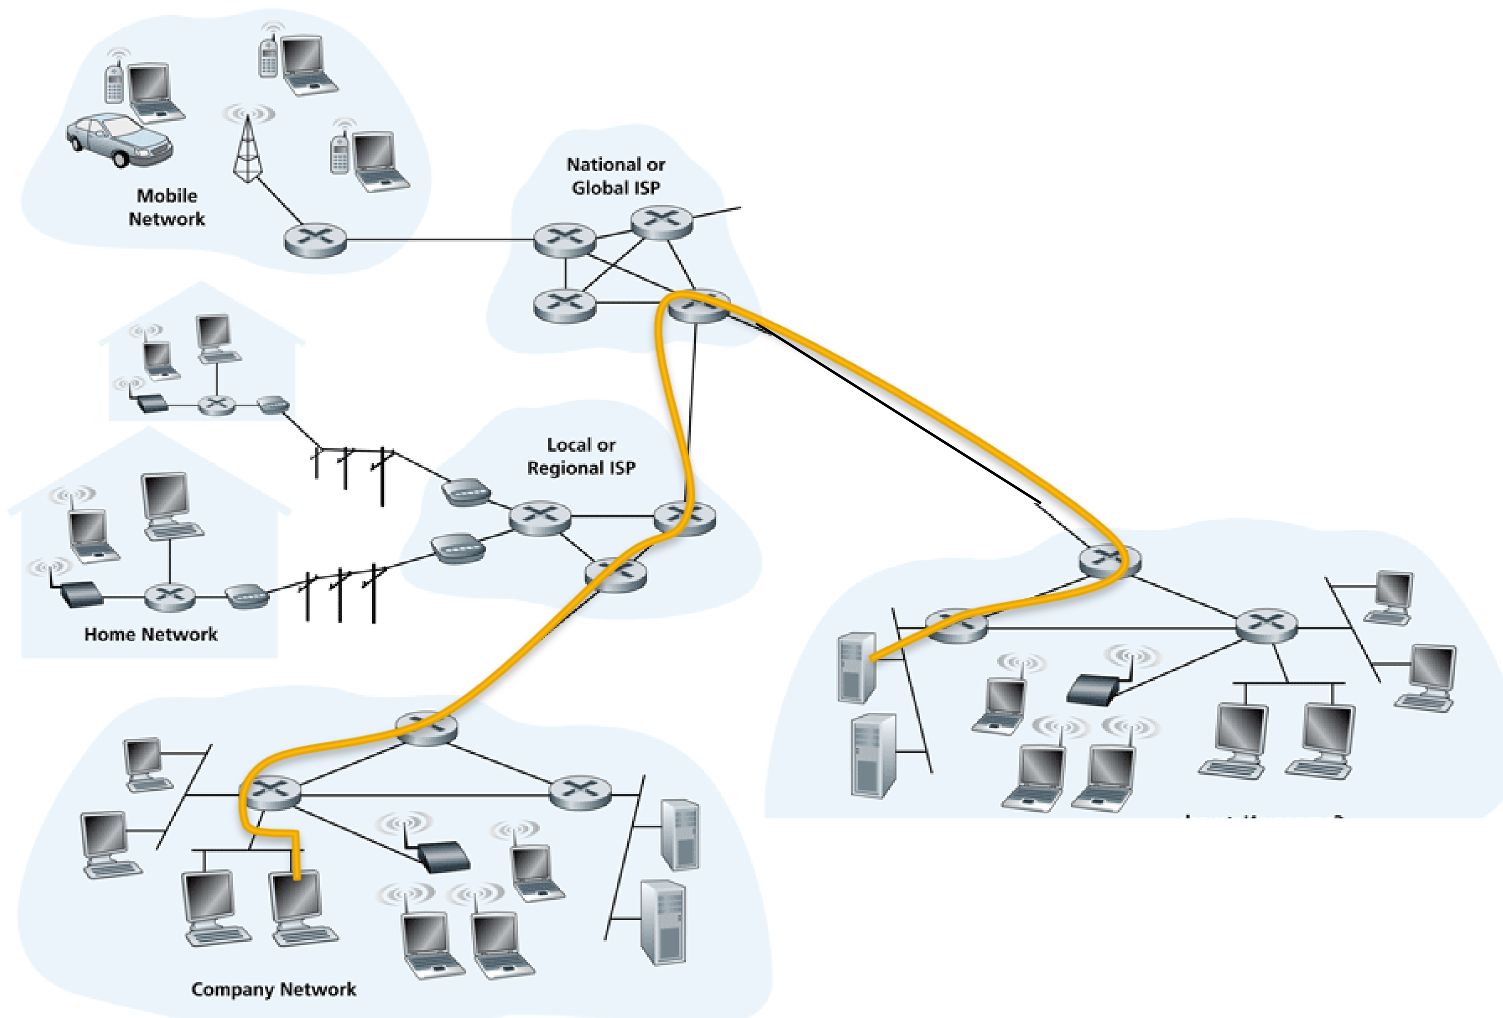

Wide-Area Services...

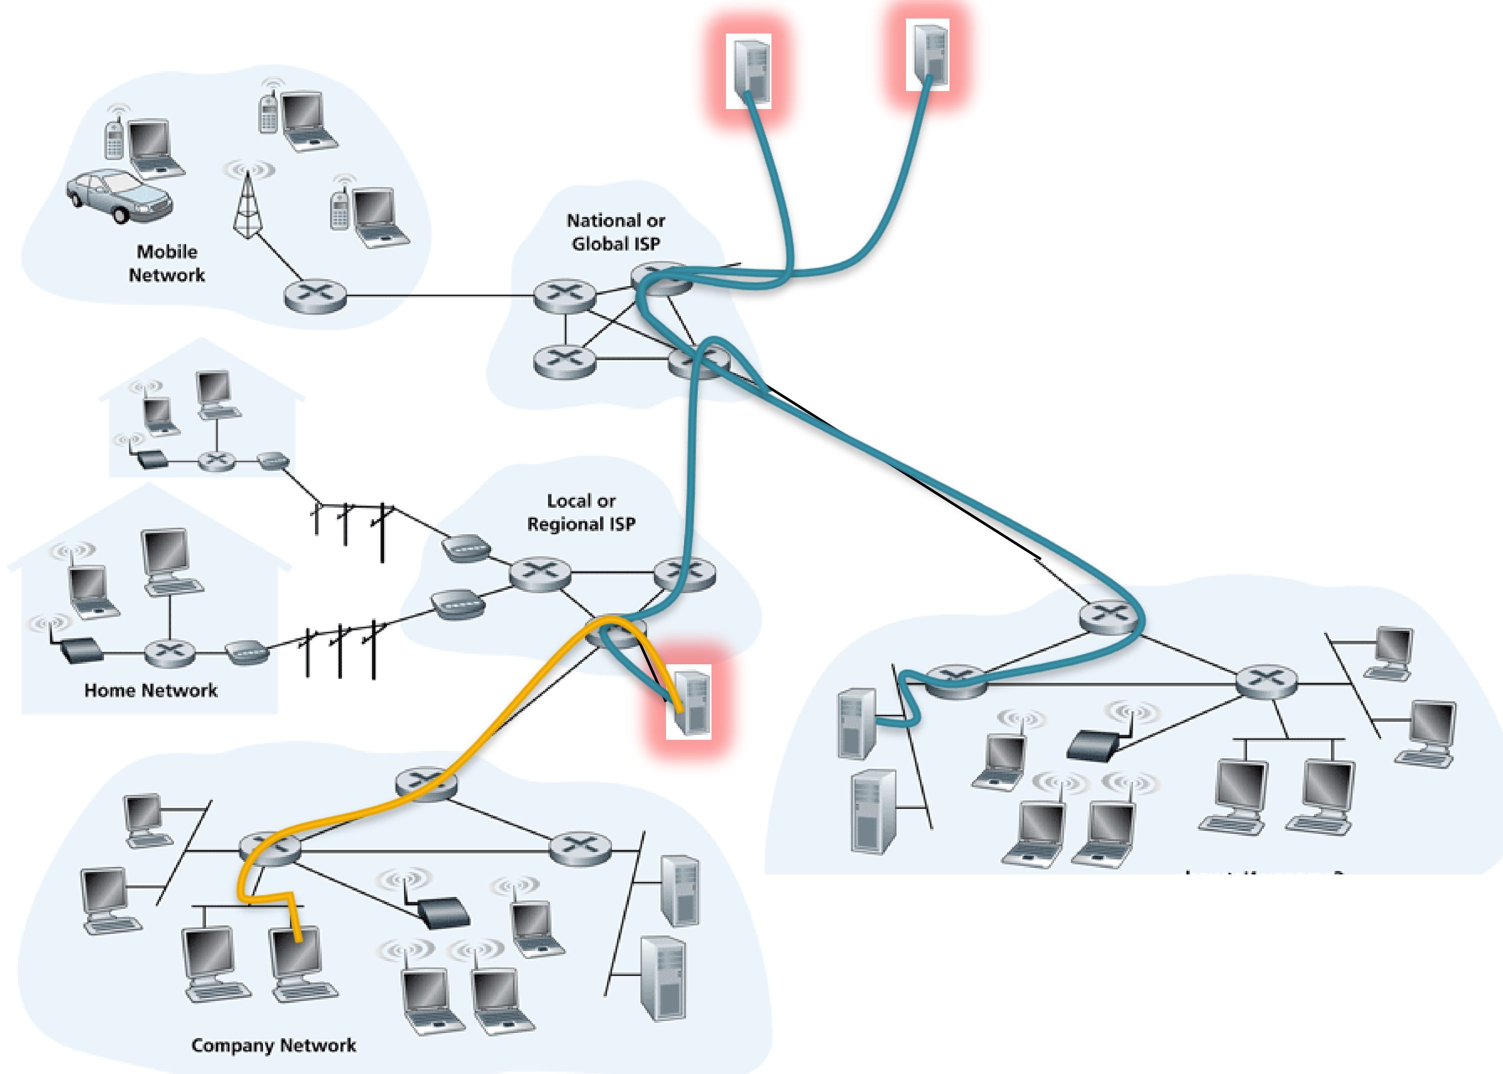

Evolution of Internet Apps

- If we look at the evolution of apps from late 90s, we can see three trends:
  - Client-server (many e-commerce, and Web-based applications)
  - Peer-to-peer (file downloads, VoIP, etc)
  - Service-oriented (Cloud computing – Amazon Elastic Computing)
- Dominant applications imposed specific requirements on the Internet infrastructure

Client-Server Applications

- Asymmetric computing
- Clients need not have persistent, globally unique addresses
  - Encouraged the proliferation of network address translators
  - Use of private address realms
- Some issues motivated by this model:
  - Performance delivered to the clients
  - Location of clients for custom services

Peer-to-peer Applications

- Symmetric computing
- End systems acting as servers and clients (downloading and uploading content)
- Middleboxes (especially firewalls and NATs) a problem for these apps
  - Applications create rendezvous points in the public network (reachable from all parties)
- Incentives for proper functioning of these apps not always clear!

Peer-to-peer Applications

- Peer-to-peer apps build "overlay" networks for resource discovery
- Two types of overlay networks:
  - Structured networks
  - Unstructured networks

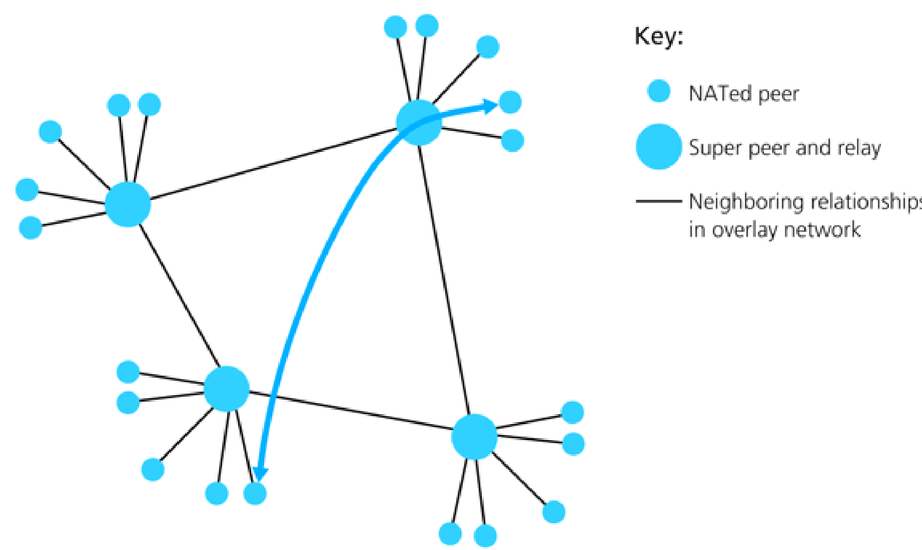

Service-Oriented Applications

- Basic idea of service-oriented computing:
  - Expose components of a large software system as web-based services
  - Use standardized web-based protocols to compose the software system
- Advantages:
  - Composition takes place at run time
  - Web infrastructure can be used to manage services (i.e., replication, placement, availability)
  - Flexibility – allows mix and match of different services
- Disadvantages:
  - Performance – not suitable for all applications

Internet Issues

- Failure – Inability to reach a given host
- Stability – oscillations in advertised routes
- Delay tolerance – network heterogeneity (e.g., satellite networks)
- Mobility – continue service while end systems hop between networks
- Attacks

Failure on the Internet

- Failure
  - Could be a naming failure (unable to resolve the name of the host to its address)
  - Could be a routing failure (unable to find a path to the host)
- Naming system failure
  - Naming system is a distributed database with a hierarchical organization
  - Uses replication to handle faults

Failure on the Internet...

- Routing failure
  - Could be a device (router) failure
  - Power-law topology actually helps here – random failures do not create much disruption on the overall function
- Could also be caused by
  - Convergence issue with the routing algorithm
  - Peering issue between two providers

"Clean Slate" Designs

- Clean Slate Research at Stanford
  - http://cleanslate.stanford.edu/
- The 100x100 Project
  - http://100x100network.org/