CENG 577 Advanced Services in Communications

Internet Architecture and Advanced Services in Converged Networks

Halûk Gümüşkaya
Fatih Üniversitesi
Bilgisayar Mühendisliği Bölümü
34500 Büyükçekmece İstanbul

Outline

- Evaluation of the Internet Architecture
- Quick Review of TCP/IP Protocol Stack
- Principles of Data Communications
- Higher Internet View and New Trends
- Business Trends
- Implications and Issues
- Summary and Conclusions

The ARPANet

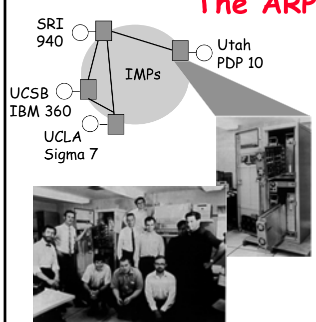

BBN team that implemented the interface message processor

- Paul Baran
  - RAND Corp, early 1960s
  - Communications networks that would survive a major enemy attack
- ARPANet: Research vehicle for "Resource Sharing Computer Networks"
  - 2 September 1969: UCLA first node on the ARPANet
  - December 1969: 4 nodes connected by phone lines

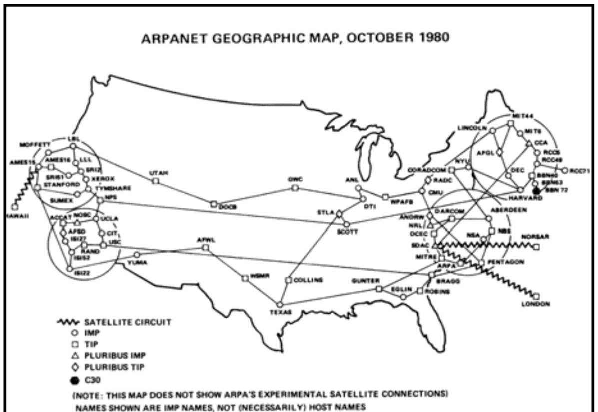

ARPANet Evolves into Internet

1965	1975	1985	1995	2005
ARPANet SATNet PRNet	TCP/IP	NFSNet	Deregulation & Commercialization WWW	ISP ASP AIP

Web Hosting
Multiple ISPs
Internet2 Backbone
Internet Exchanges

→

Application Hosting
ASP: Application Service Provider
AIP: Application Infrastructure Provider (e-commerce tookit, etc.)

NSFNET T1 Network 1991

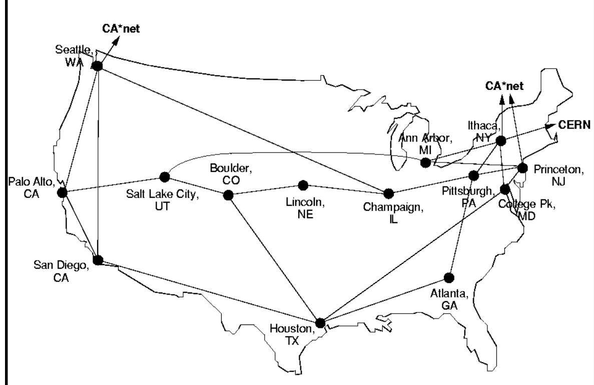

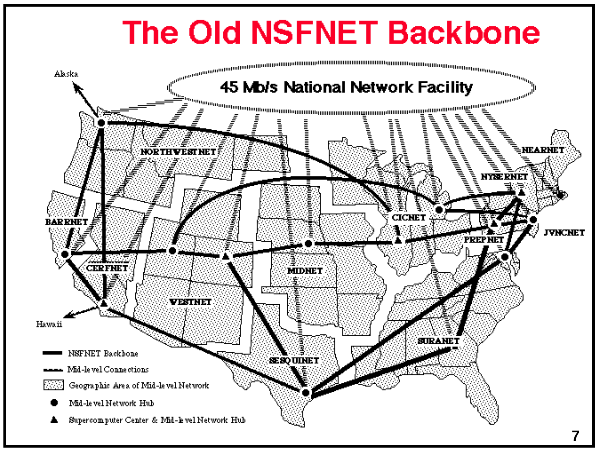
The Old NSFNET Backbone
45 Mb/s National Network Facility
Alaska
Hawaii
NORTHWESTNET
BARRNET
CERFNET
WESTNET
MIDNET
CICNET
NEARNET
NYSERNET
JVNCNET
PREPNET
SURANET
SESQUINET
NSFNET Backbone
Mid-level Connections
Geographic Area of Mid-level Network
Mid-level Network Hub
Supercomputer Center & Mid-level Network Hub
7

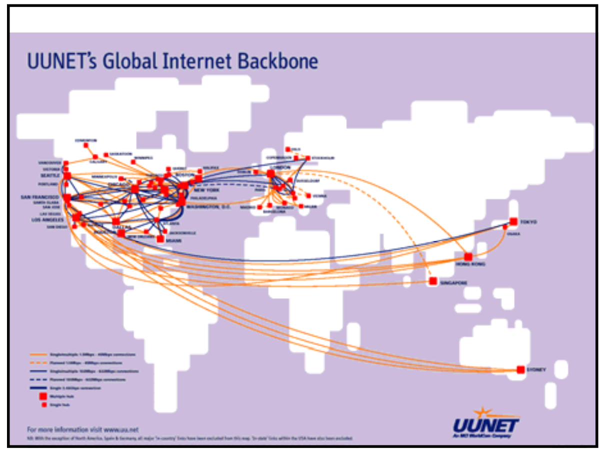
UUNET's Global Internet Backbone
UUNET

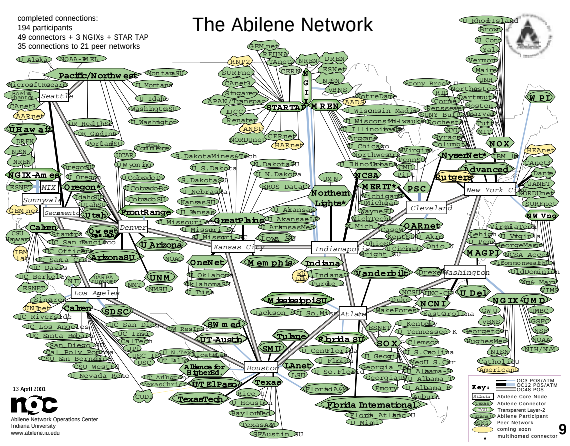
The Abilene Network
completed connections:
194 participants
49 connectors + 3 NGIXs + STAR TAP
35 connections to 21 peer networks
13 April 2001
Abilene Network Operations Center
Indiana University
www.abilene.iu.edu
9

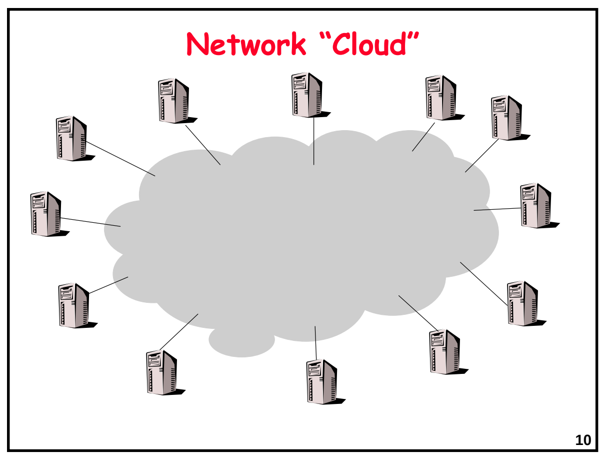
Network "Cloud"
10

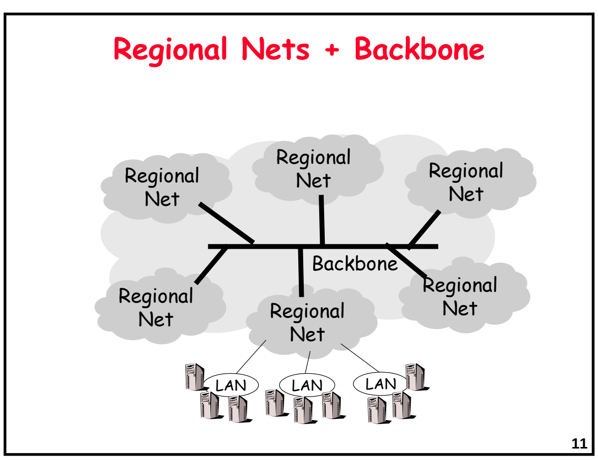
Regional Nets + Backbone
Regional Net
Regional Net
Regional Net
Backbone
Regional Net
Regional Net
Regional Net
LAN
LAN
LAN
11

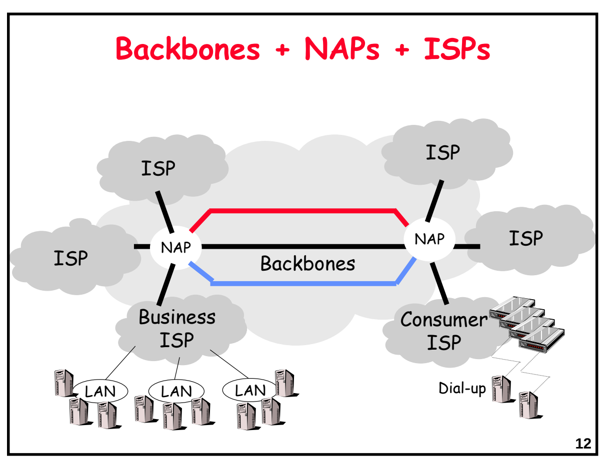
Backbones + NAPs + ISPs
ISP
ISP
ISP
NAP
Backbones
NAP
ISP
Business ISP
Consumer ISP
LAN
LAN
LAN
Dial-up
12

Core Networks + Access Networks

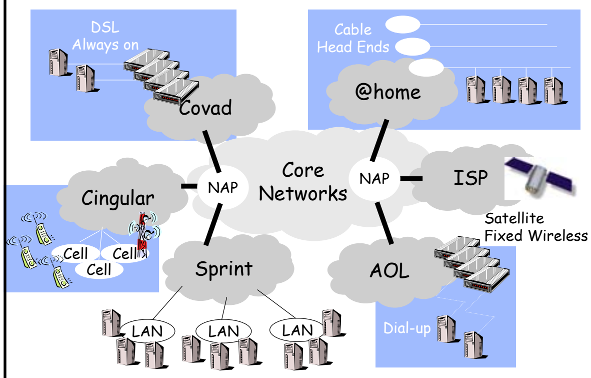

13

Computers Inside the Core

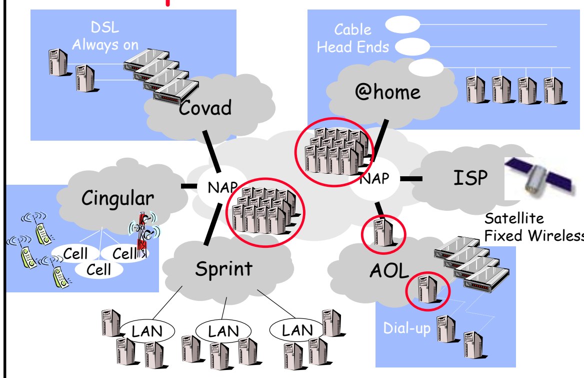

14

Connectivity and Processing

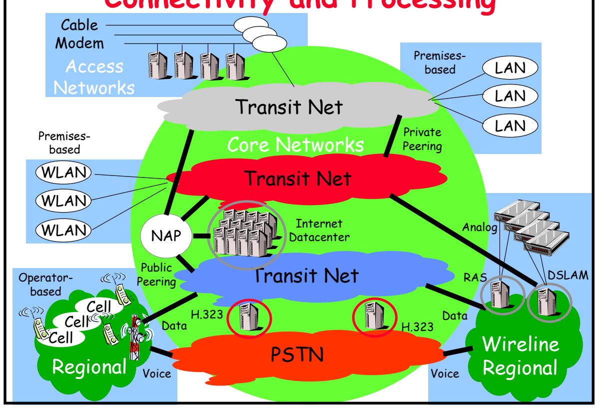

15

Outline

- Evaluation of the Internet Architecture
- Quick Review of TCP/IP Protocol Stack
- Principles of Data Communications
- Higher Internet View and New Trends
- Business Trends
- Implications and Issues
- Summary and Conclusions

16

What is the Internet?
"It's the TCP/IP Protocol Stack"

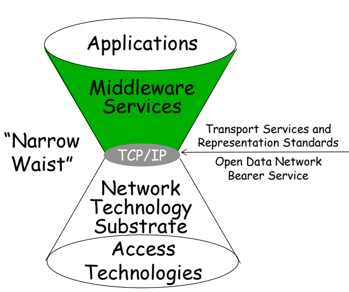

- Applications
  - Web
  - Email
  - Video/Audio
- TCP/IP
- Access Technologies
  - Ethernet (LAN)
  - Wireless (LMDS, WLAN, Cellular)
  - Cable
  - ADSL
  - Satellite

Where is the next "narrow waist"?

17

OSI Reference Model

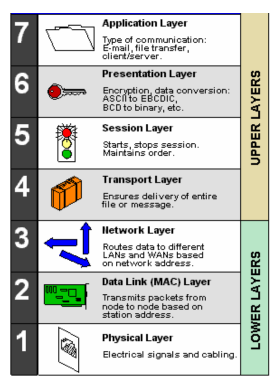

18

OSI and TCP/IP Reference Models

OSI Model	TCP/IP	TCP/IP	Security Protocols
Application	Application	Application	S-MIME S-HTTP SET ...
Presentation			
Session			
Transport	Transport (TCP/UDP)	Transport (TCP/UDP)	SOCKS, SSL, TSL
Network	IP	IP	IPSEC (AH, ESP) Packet Filtering Tunneling
Data Link	Data Link	Network Access	Link Encryption
Physical	Physical		

TCP/IP (Internet) Protocol Stack

- **Application:** supporting network applications
  - HTTP, FTP, SMTP, POP3, SNMP, …
- **Transport:** host-host data transfer
  - TCP, UDP
- **Network:** routing of datagrams from source to destination
  - IP, routing protocols
- **Link:** data transfer between neighboring network elements
  - PPP, Ethernet, 802.11, …
- **Physical:** bits "on the wire"

Application
Transport (TCP/UDP)
IP
Data Link
Physical

TCP/IP (Internet) Protocol Stack

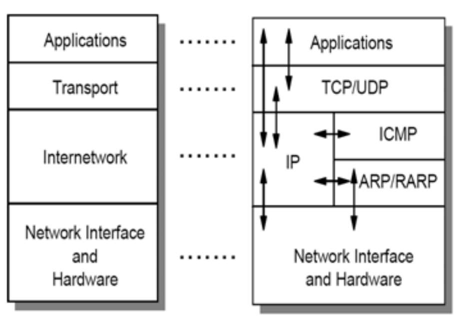

Each layer represents a package of functions

Detailed TCP/IP (Internet) Protocol Stack

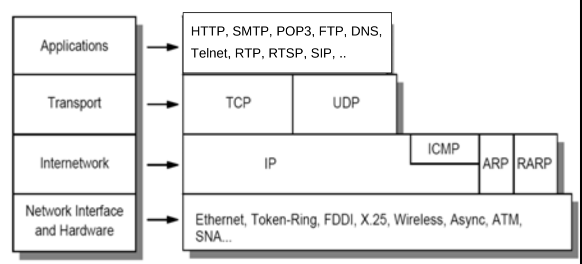

Do you want more details?

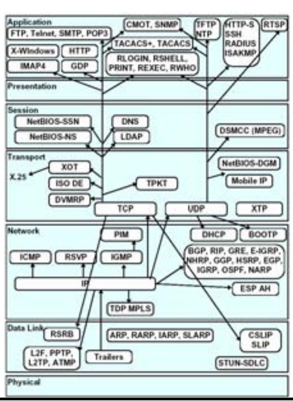

Outline

- Evaluation of the Internet Architecture
- Quick Review of TCP/IP Protocol Stack
- Principles of Data Communications
- Higher Internet View and New Trends
- Business Trends
- Implications and Issues
- Summary and Conclusions

Data Traveling in Protocol Stack

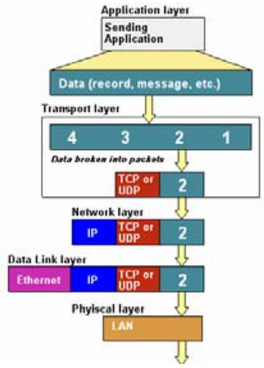

25

Data Traveling in Protocol Stack

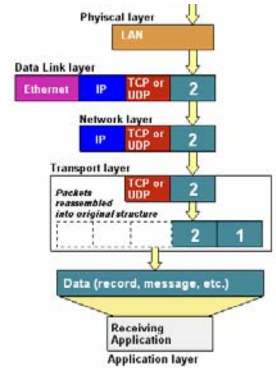

26

Connection in TCP/IP

A connection between two machines in TCP/IP is defined by:

- Transport layer protocol (TCP or UDP)
- IP address of local machine
- Port number used on the local machine
- IP address of remote machine
- Port number used on the remote machine

27

Transport-Level Protocols: Ports

- TCP and UDP protocols use **ports** to map incoming data to a particular process running on a computer.
- IP Datagram identifies the **host** and the **port** that it's destined for.
- The **computer is identified** by its **32-bit IP address**, which IP uses to deliver data to the right computer on the network.
- **Ports are identified by** a 16-bit integer number, ranging from 0 to 65535, which TCP/UDP use to deliver the data to the right application.

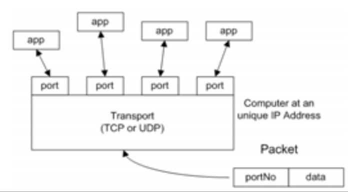

28

Well-known Ports

- Port numbers between 0 and 1023 are restricted (**well-known ports**) -- they are reserved for use by well-known services such as HTTP and FTP and other system services.
- Your applications should not attempt to bind to these ports.

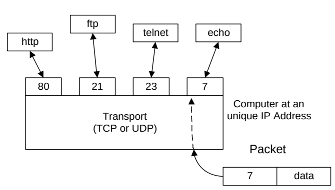

29

Well-Known TCP/IP Services and Port Assignments

Protocol	*Port*	*Encoding*	*Purpose*
Echo	7	TCP/UDP	Test protocol used to verify that two machines are able to connect by having one echo back the other's input.
Discard	9	TCP/UDP	Less useful test protocol that ignores all data received by the server.
Daytime	13	TCP/UDP	Provides an ASCII representation of the current time on the server.
ftp-data	20	TCP	FTP uses two well-known ports. This port is used to transfer files.
FTP	21	TCP	This port is used to sent ftp commands like, "put" and "get".
TELNET	23	TCP	A protocol used for interactive, remote command-line sessions.
SMTP	25	TCP	"Simple Mail Transfer Protocol" is used to send email between machines.
Time	37	TCP/UDP	A time server returns the number of seconds that have elapsed on the host machine since midnight, January 1, 1900, as a four-byte, signed, big-endian integer.
Whois	43	TCP	Simple directory service for Internet network administrators.
Finger	79	TCP	It gets information about a user or users.
HTTP	80	TCP	Hyper Text Protocol is the underlying protocol of the World Wide Web.
POP3	110	TCP	Post Office Protocol version 3 is a protocol for the transfer of accumulated email form the host to sporadically connected clients.
NNTP	119	TCP	Usenet news transfer. More formally known as the "Network News Transfer Protocol".
SNMP	161/162	UDP	Simple Network Managment Protocol is used in management of TCP/IP.
RMI Reg.	1099	TCP	The RMI Registry is a registry service for Java remote objects.
Servlets	8080	TCP	Java Server API and Servlets is a web server from Sun that runs on port 8080 by default, not port 80.

30

Transport-Level Protocols: TCP

- TCP is a **reliable** and **connection-oriented** communication protocol on top of the unreliable, unsequenced functionality of IP.

Reliable:

- TCP provides extensive **error-checking** capabilities.
- TCP provides **reliable stream delivery**. This reliable stream delivery ensures that the data arrives in the same sequence and state in which it was sent.

Connection-oriented:

- The TCP system relies on a **virtual circuit** that is established between the requesting machine and its target.
- This circuit is opened via a 3-part process, often referred to as the **3-part handshake**.

31

Transport-Level Protocols: TCP

- Because of the reliable and sequenced nature of TCP sockets, they often are called **stream sockets**; you can read and write data in continuous streams of bytes without worrying about packets, headers, and so on.
- TCP is the chief protocol employed on the Internet.
- It facilitates such mission-critical tasks as file transfers and remote sessions.
- Stream socket functionality in Java is provided by the classes **java.net.ServerSocket** and **java.net.Socket**.

32

Transport-Level Protocols: UDP

- UDP is an **unreliable** and **connectionless** communication protocol.
- Datagram-based communication.
- Datagram packets are prepared by the applications.
- IP Address + Port Number are put into datagram.
- UDP-based communication is like sending letters to a post office.
- Not reliable but fast compared to TCP.
- Datagram socket functionality in Java is provided by the classes **java.net.DatagramSocket** and **java.net.DatagramPacket**.

33

TCP-Based Communication Steps

- A **server** application opens a socket to establish a connection with another application (client) by binding a socket to a port number. (registering the application with the system to receive all data destined for that port.)
- Server: TCP Socket = Port Number (Well-known)
- Client: TCP Socket = IP Address + Port Number (server's port)
- When a **client** makes a request from the server's port, input and/or output streams are created on the socket depending on the protocol used between the server and the client.
- No two applications can bind to the same port: Attempts to bind to a port that is already in use will fail.
- Stream based (like a phone call)
- Uses 3-way handshake, reliable but slow (compared to UDP)

34

Iterative and Concurrent TCP Servers

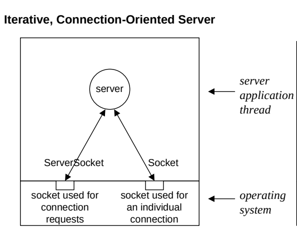

A server implementation that processes one request at a time.

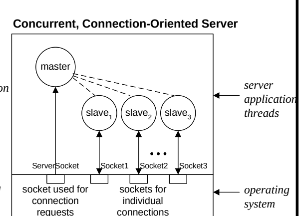

Concurrent server handles multiple requests at one time.

35

UDP Client or Server

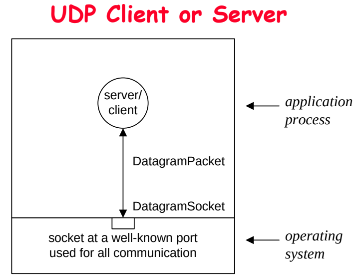

The same socket is used to send data and to listen for incoming connections. Applications handle the client/server functionality.

36

Outline

- Evaluation of the Internet Architecture
- Quick Review of TCP/IP Protocol Stack
- Principles of Data Communications
- Higher Internet View and New Trends
- Business Trends
- Implications and Issues
- Summary and Conclusions

37

Internet Applications

- **Variations on three themes**
  - distinguish protocol vs. application behavior
- **Messaging**
  - datagram model → no direct confirmation of final receipt
  - email (optional confirmation now) and IM
  - emphasis on interoperation (SMS, pagers, …)
  - delays measured in minutes
- **Retrieval & query (request/response)**
  - "client-server"
  - ftp, HTTP
  - RPC (Sun RPC, DCE, DCOM, Corba, XML-RPC, SOAP)
  - emphasis on fast & reliable transmission
  - delays measured in seconds

38

Internet applications, cont'd

- **Continuous Media**
  - generation rate ~ delivery rate ~ rendering rate
  - audio, video, measurements, control
    - » Internet telephony
    - » Multimedia conferencing
  - related: streaming media →slightly longer timescales for rate matching
    - » video-on-demand
  - emphasis is on *timely* and low-loss delivery → *real-time*
  - delays measured in milliseconds
  - focus of this course

39

Internet protocols

- Protocols support these applications:
  - data delivery
    - » HTTP, ftp data part, SMTP, IMAP, POP, NFS, SMB, RTP
  - identifier mapping (id → id, id → data)
    - » ARP, DNS, LDAP
  - configuration (= specialized version of identifier → data)
    - » DHCP, ACAP, SLP, NETCONF, SNMP
  - control and setup
    - » RTSP, SIP, ftp control, RSVP, SNMP, BGP and routing protocols
- May be integrated into one protocol or general service function ("middleware"?)

40

Networking is getting into middle years

	idea	current
IP	1969, 1980?	1981
TCP	1974	1981
telnet	1969	1983
ftp	1980	1985

41

Standardization

- Really two facets of standardization:
  1. public, interoperable description of protocol, but possibly many (Tanenbaum)
  2. reduction to 1-3 common technologies
     - LAN: Arcnet, tokenring, ATM, FDDI, DQDB, … → Ethernet
     - WAN: IP, X.25, OSI → IP
- Have reached phase 2 in most cases, with RPC (SOAP) and presentation layer (XML) most recent 'conversions'

42

Technologies at ~30 years

- Other technologies at similar maturity level:
  - air planes: 1903 – 1938 (Stratoliner)
  - cars: 1876 – 1908 (Model T)
  - analog telephones: 1876 – 1915 (transcontinental telephone)
  - railroad: 1800s -- ?

43

Observations on progress

- 1960s: military → professional → consumer
  - now, often reversed
- Oscillate: convergence → divergence
  - continued convergence clearly at physical layer
  - niches larger → support separate networks
- Communications technologies rarely disappear (as long as operational cost is low):
  - exceptions:
    - » telex, telegram, semaphores → fax, email
    - » X.25 + OSI, X.400 → IP, SMTP
  - analog cell phones

44

History of networking

- History of networking = non-network applications migrate
  - postal & intracompany mail, fax → email, IM
  - broadcast: TV, radio
  - interactive voice/video communication → VoIP
  - information access → web, P2P
  - disk access → iSCSI, Fiberchannel-over-IP

45

Network evolution

- Only three modes, now thoroughly explored:
  - packet/cell-based
  - message-based (application data units)
  - session-based (circuits)
- Replace specialized networks
  - left to do: embedded systems
    - » need cost(CPU + network) < $10
    - » cars
    - » industrial (manufacturing) control
    - » commercial buildings (lighting, HVAC, security; now LONworks)
    - » remote controls, light switches
    - » keys replaced by biometrics

46

New applications

- New bandwidth-intensive applications
  - Reality-based networking
  - (security) cameras
- Distributed games often require only low-bandwidth control information
  - current game traffic ~ VoIP
- Computation vs. storage vs. communications
  - communications cost has decreased less rapidly than storage costs

47

Security challenges

- DOS, security attacks → permissions-based communications
  - only allow modest rates without asking
  - effectively, back to circuit-switched
- Higher-level security services → more application-layer access via gateways, proxies, …
- User identity
  - problem is not availability, but rather over-abundance

48

Scaling

- Scaling is only backbone problem
- Depends on network evolution:
  - continuing addition of AS to flat space → deep trouble
  - additional hierarchy

49

Quality of Service (QoS)

- QoS is meaningless to users
- care about service availability → reliability
- as more and more value depends on network services, can't afford random downtimes

50

Textbook Internet vs. real Internet

end-to-end (application only in 2 places)	middle boxes (proxies, ALGs, …)
permanent interface identifier (IP address)	time-varying (DHCP)
globally unique and routable	network address translation (NAT)
multitude of L2 protocols (ATM, ARCnet, Ethernet, FDDI, modems, …)	dominance of Ethernet, but also L2's not designed for networks (1394 Firewire, Fibre Channel, MPEG2, …)

51

Textbook Internet vs. real Internet

mostly trusted end users	hackers, spammers, con artists, pornographers, …
small number of manufacturers, making expensive boxes	Linksys, Dlink, Netgear, …, available at Radio Shack
technical users, excited about new technology	grandma, frustrated if email doesn't work
4 layers (link, network, transport, application)	layer splits
transparent network	firewalls, L7 filters, "transparent proxies"

52

Internet architecture documents (readings)

- http://www.ietf.org/rfc/rfcXXXX.txt
- RFC 1287
- RFC 2101
- RFC 2775
- RFC 3234

53

The Internet Protocol Hourglass (Deering)

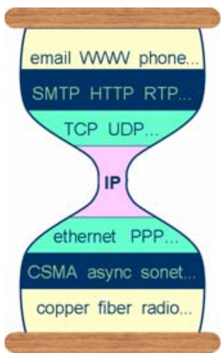

54

Why the hourglass architecture?

- **Why an internet layer?**
  - make a bigger network
  - global addressing
  - virtualize network to isolate end-to-end protocols from network details/changes
- **Why a *single* internet protocol?**
  - maximize interoperability
  - minimize number of service interfaces
- **Why a *narrow* internet protocol?**
  - assumes least common network functionality to maximize number of usable networks

Deering, 1998

55

Putting on Weight

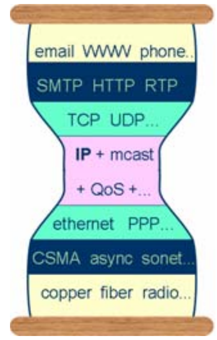

- requires more functionality from underlying networks

56

Mid-Life Crisis

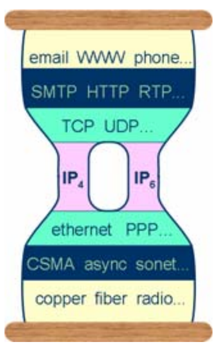

- doubles number of service interfaces
- requires changes above & below
- major interoperability issues

57

IPv6

- Initial motivation: 32-bit address space soon to be completely allocated.
- Additional motivation:
  - header format helps speed processing/forwarding
  - header changes to facilitate QoS

  IPv6 datagram format:
  - fixed-length 40 byte header
  - no fragmentation allowed

58

IPv6 Header (Cont)

Priority: identify priority among datagrams in flow
Flow Label: identify datagrams in same "flow."
(concept of "flow" not well defined).
Next header: identify upper layer protocol for data

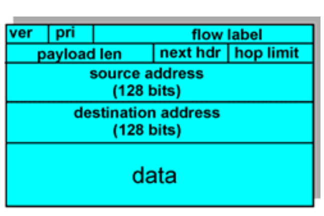

59

Other Changes from IPv4

- *Checksum*: removed entirely to reduce processing time at each hop
- *Options:* allowed, but outside of header, indicated by "Next Header" field
- *ICMPv6:* new version of ICMP
  - additional message types, e.g. "Packet Too Big"
  - multicast group management functions

60

Transition From IPv4 To IPv6

- Not all routers can be upgraded simultaneous
  - no "flag days"
  - How will the network operate with mixed IPv4 and IPv6 routers?
- *Tunneling:* IPv6 carried as payload in IPv4 datagram among IPv4 routers

61

Tunneling

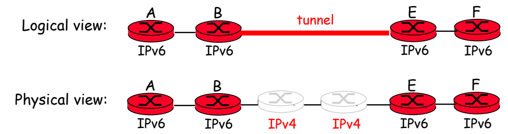

62

Tunneling

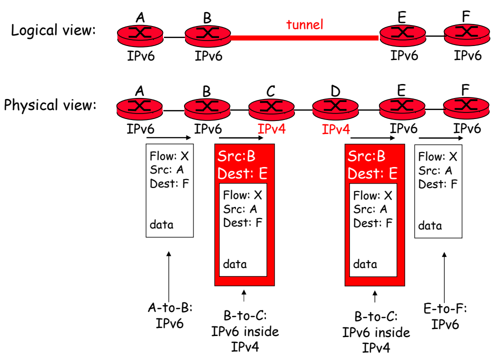

63

Layer splitting

- Traditionally, L2 (link), L3 (network = IP), L4 (transport = TCP), L7 (applications)
- Layer 2: Ethernet → PPPoE (DSL)
- Layer 2.5: MPLS, L2TP
- Layer 3: tunneling (e.g., GPRS)
- Layer 4: UDP + RTP
- Layer 7: HTTP + real application

64

Layer violations

- Layers offer abstraction → avoid "Internet closed for renovation"
- Cost of information hiding
- Cost of duplication of information when nothing changes
  - fundamental design choice of Internet = difference between circuit and datagram-oriented networks
- Assumption: packets are large and getting larger
  - wrong for games and audio
- Cost prohibitive on wireless networks
  - will see: 10 bytes of payloads, 40 bytes of packet header
  - header compression → compress into state index on one link

65

Internet acquires presentation layer

- All learn about OSI 7-layer model
- OSI: ASN.1 as common rendering of application data structures
  - used in LDAP and SNMP (and H.323)
- Internet never really had presentation layer
  - approximations: common encoding (TLV, RFC 822 styles)
- Now, XML as the design choice by default

66

Internet acquires session layer

- Originally, meant for data sessions
- Example (not explicit): ftp control connection
- Now, separate data delivery from session setup
  - address and application configuration
  - deal with mobility
  - will see as RTSP, SIP and H.323

67

Outline

- Evaluation of the Internet Architecture
- Quick Review of TCP/IP Protocol Stack
- Principles of Data Communications
- Higher Internet View and New Trends
- **Business Trends**
- Implications and Issues
- Summary and Conclusions

68

Layerized Internet Service Business Model

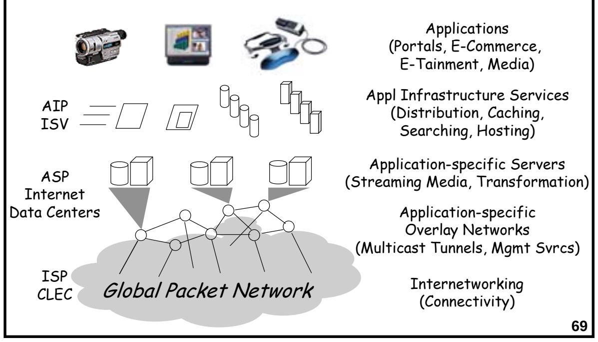

69

A New Kind of Internet

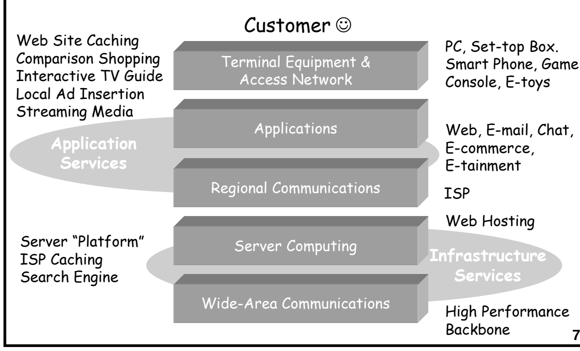

70

Open vs. Closed Access to Services

Covad DSL	Time/Warner Roadrunner AOL Dial-up	AT&T Cable	Access	Cable, DSL, MMDS, LMDS, Satellite
CNCX	AOL	@Home	Local Network Management	ISP
Williams	AOL	@Home	Routing & Distribution	Backbone Provider
Web	AOL/Netscape Time/Warner	Excite	Content	Portal Web Sites

- Closed end-to-end pipe: optimized performance
- But companies developing compelling infrastructure technology that any content provider or ISP can adopt
- Closed system can't benefit from these

71

From Network Management to *Service* Management

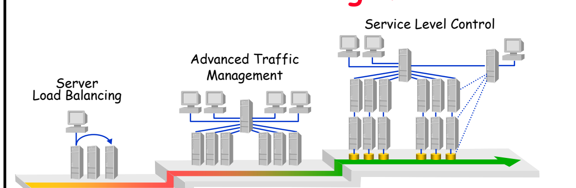

- Server and site availability
- Balanced server and site load

- Rapid change
- Network and application flexibility
- Scalability
- Complex site administration

- Rapid problem diagnosis/ isolation
- Service level measurement
- Multi-tier resource monitoring

- Preferential Services
- Resource Provisioning
- Self-tuning
- Problem prevention

Chris Morino, Resonate

72

Service Reliability is Critical

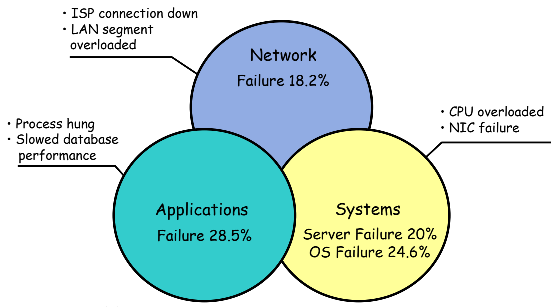

Source: IDC

Administration 8.7%

Chris Morino, Resonate 73

Competition vs. Cooperation

- Internet Service Providers: Competition
  - Peering for packet transport: BGP protocol
  - Charging based on traffic volumes

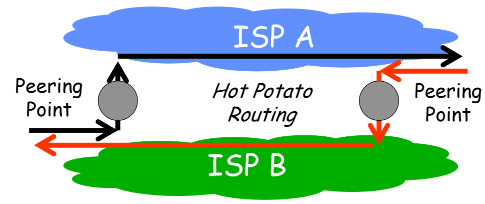

74

Mobile Internet Might Be Different Than Wired Internet

- Wireless is a smarter pipe
  - Location-awareness
  - UI dictates need for personalization, mediation
- Clear billing authority: it's the access provider
  - People actually do pay for transport
  - Reverse billing allows content provider to charge for service
- Peering as a necessity
  - Operators provide local service
  - Roaming agreements provide basis for service peering
  - Well understood arrangements for settlements
  - New economies driving towards shared network deployment
- Person-to-Person communications *is* a killer app
- Microsoft's non-monopoly

75

Cooperation and Peering

- 3G Spectrum Auctions: 150 billion ECU; Capital outlays may match spectrum expenses, all before first revenue
- New business models in Mobile Networks
  - Compelling services make the difference
  - Collaborate on deployment of physical network
  - Compete on provisioning of services
- Peering For More Than Connectivity
  - Horizontal architecture of services on top of networks
  - Virtual Home Environments
  - Relationships between operators, billing agents, service providers

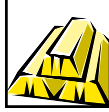

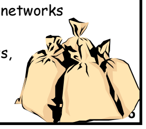

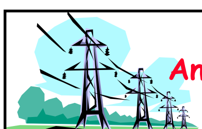

Any Way to Build a Network?

- Partitioning of frequencies independent of actual subscriber density
  - Successful operator oversubscribe resources, while less popular providers retain excess capacity
  - Different flavor of roaming: among collocated/competing service providing
- Duplicate antenna sites
  - Serious problem given community resistance
- Redundant backhaul networks
  - Limited economies of scale

77

The Case for Horizontal Architectures

"The new rules for success will be to provide one part of the puzzle and to cooperate with other suppliers to create the complete solutions that customers require. ... [V]ertical integration breaks down when innovation speeds up. The big telecoms firms that will win back investor confidence soonest will be those with the courage to rip apart their monolithic structure along functional layers, to swap size for speed and to embrace rather than fear disruptive technologies."
The Economist Magazine, 16 December 2000

78

Feasible Alternative: Horizontal *Competition* vs. Vertical *Integration*

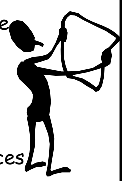

- Service Operators "own" the customer, provide "brand", issue/collect the bills
- Independent Backhaul Operators
- Independent Antenna Site Operators
- Independent Owners of the Spectrum
- Microscale auctions/leases of network resources
- Emerging concept of Virtual Operators

79

Business as Usual: Vertical Integration

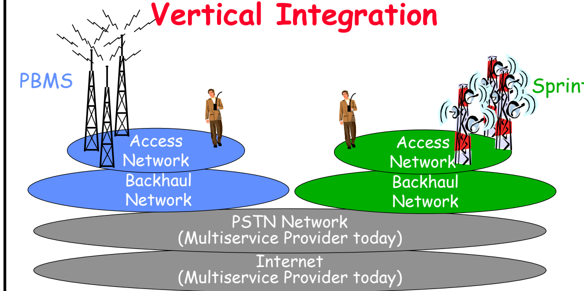

- Each operator owns own frequencies, cell sites, backhaul network

80

Business Unusual: Horizontal Competition

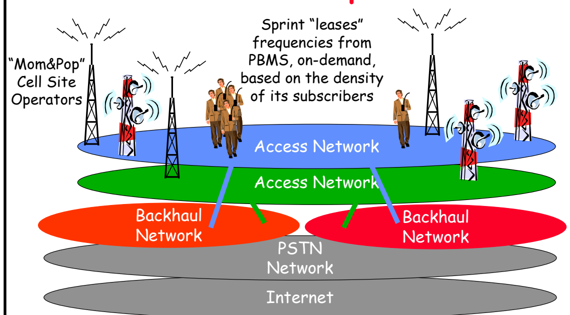

81

Virtual Operator

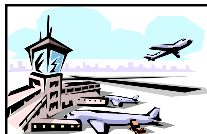


- MVNO: Virgin Mobile and One2One in UK
  - Distinguish based on marketing and billing plan innovations
  - VM competes for subscribers but uses One2One's network
- "Operators without subscribers": local premises deploy own access infrastructure
  - Better coverage/more rapid build out of network
  - Deployments in airports, hotels, conference centers, office buildings, campuses, …
- Overlay service provider (e.g., PBMS) vs. organizational service provider (e.g., UCB IS&T)
  - Single bill/settle with service participants
  - Operator Wireless LAN
- Support for ensemble devices
  - Cell Phone + Wall Camera & Display

82

Outline

- Evaluation of the Internet Architecture
- Quick Review of TCP/IP Protocol Stack
- Principles of Data Communications
- Higher Internet View and New Trends
- Business Trends
- **Implications and Issues**
- Summary and Conclusions

83

What will be the Next Generation of Driving Applications?

- Location-aware/context-aware information delivery and presentation
  - Extends UniIn-Box: loc-based, exploits calendar info
  - Mediation to translate formats
- IP Telephony, Packet VoD, Teleconferencing
  - Streaming media, multicast-based
  - Bandwidth, latency, jitter, lose rate constraints
  - Clearinghouse provisioning
- Event Delivery for Distributed Applications
  - Performance/reliability constrained messaging
  - Management of Content Delivery Networks, Distributed Service architecture?
- Interactive Games? Distributed Storage (OceanStore)? Telemetry?

84

What Will Be the Next Generation Operational Environment?

- Virtual Operators/Service Provider (VOSP)
  - Provide service to end users with no server/network infrastructure of own
  - Independent "Path" providers (e.g., ISPs) and Server providers (e.g., Internet Data Centers)
  - Many-to-many relationship between VOSP and Path/Server Providers
- Confederated Service Provider
  - Service-level peering: sharing of paths and servers to deploy end-to-end service with performance and reliability constraints
- Note: Akamai runs "the world's largest service network" without owning a network!

85

Alternative Operational Environments

- Confederation Model
  - Providers share (limited) information about topology, server location, path performance
  - Cooperatively collect internal information and share
- Overlay Model
  - Reverse-engineer topology and intra-cloud performance
  - Collection done by brokers outside of the cloud
- SLAs, Verification, Maintenance of Trust Relationships different in the two models
- Is there an operational/performance advantage to the Confederation Model?

86

Open Issues/Questions

- Traditional Overlay Networks
  - Server ("Application Level Router") Placement
    - » For scaling, reliability, load balancing, latency
    - » Where? Network topology discovery: WAN Core, Metro/Regional, Access Networks
  - Choice of Inter-Server "Paths"
    - » For server-to-server latency/bandwidth/loss rate
    - » Predictable/verifiable network performance (intra-ISP SLA)
  - Redirection Mechanisms
    - » Random, round-robin, load-informed redirection
    - » Net vs. server as bottleneck

87

Open Issues/Questions

- Performance-constrained Service Placement
  - Separation of Service, Server, Service Path
    - » Assume "Server Centers" known, can be "discovered" (how does OceanStore deal with this?), or register with a Service Placement Service (SPS)
    - » How is Service named, described, performance constraints expressed, and registered?
    - » How is app/service-specific performance measured and made known to Service Placement Service?
  - Brokering between Server Centers and Service Creator, Path Provider and Service Creator
- If core network bandwidth becomes infinite and "free", does it matter where services are placed?
  - Latency reduction vs. economies of centralized management

88

Emerging Reference Architecture

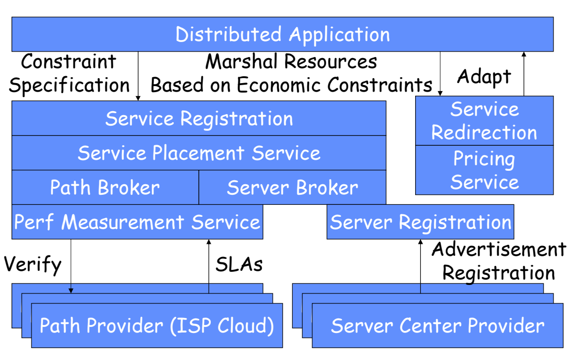

89

Methodological Framework

- Problem: performing scaled, wide-area networking studies in the current Internet environment
- Possible Solution: Wide-area Network Emulation
  - Virtual WAN (VWAN) on Large-scale Multicomputer Testbeds
  - Build operational model on top of VWAN
    - » Traffic generation and measurement infrastructure
    - » Build Confederation and Overlay operational models
    - » What part of mechanisms for measurement, negotiation, registration, redirection, etc. the same and which are different?

90

Outline

- Evaluation of the Internet Architecture
- Quick Review of TCP/IP Protocol Stack
- Principles of Data Communications
- Higher Internet View and New Trends
- Business Trends
- Implications and Issues
- Summary and Conclusions

91

Connectivity and Processing

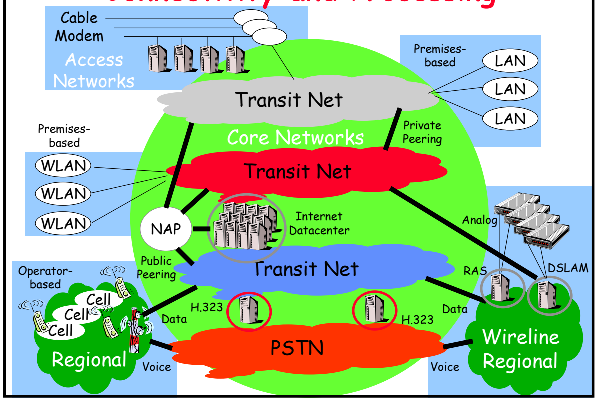

Challenges for Converged Networks

- Services spanning access networks, to achieve high performance and manage diversity of end devices
- Not about specific Information Appliances
- Builds on the New Internet: multiple application-specific "overlay" networks, with new kinds of *service-level* peering
- Pervasive support for services within "intelligent" networks
  - Automatic replication
  - Document routing to caches
  - Compression & mirroring
  - Data transformation

93

Managing Edge Versus Core Services

- Wide-area bandwidth efficiency
- Increasing b/w over access networks, but impedance mismatch between core and access nets
- Fast response time (and more predictable)
- Opportunity to untegrate localized content
- Associated with client (actually ISP), not server
- Examples:
  - Caching: exploits response time, b/w efficiency, high local b/w
  - Filtering: form of local content transformation
  - Internet TV: b/w efficiency, high local b/w, predictable response
  - Transformation: adapt content for end user/diverse access devices
  - Software Rental: sxploits high local b/w
  - Games, chat rooms, ….

94

Yielding a New Research Agenda

- New Definition of "Quality of Service"
  - Perceived quality depends on services in the network
  - Manage caches, redirection, NOT bandwidth
  - Enable incorporation of localized content
- Bandwidth Issues
  - Tier 1 ISP backbones rapidly moving towards OC 192 (9.6 gbs!)
  - Better interconnection: hops across ASs decreasing over time
  - Emerging broadband access networks: cable, DSL, …
  - End-to-end latency/server load dominate performance
- Supporting Old Services in the New Internet
  - IP Multicast, DNS, …
  - Rethinking the End-to-End Principle
  - Service/content-level peering, just like routing-level peering
  - Secure end-to-end connection compatible with service model?

95

Acknowledgements

- Slides have been based in-part upon original slides/notes of a number of people, including:
- Randy H. Katz, University of California, Berkeley
- Henning Schulzrinne, Columbia University
- Jim Kurose, University of Massachusetts

96