

Operating Systems Design

13. Client-Server Networking

Paul Krzyzanowski
pzk@cs.rutgers.edu

Modes of connection

Circuit-switched

- dedicated path
- guaranteed (fixed) bandwidth
- [almost] constant latency

Packet-switched

- shared connection
- data is broken into chunks called packets
- each packet contains destination address
- available bandwidth \leq channel capacity
- variable latency

What's in the data?

For effective communication

- same language, same conventions

For computers:

- electrical encoding of data
- where is the start of the packet?
- which bits contain the length?
- is there a checksum? where is it?
how is it computed?
- what is the format of an address?
- byte ordering

Protocols

These instructions and conventions
are known as **protocols**

Protocols

Exist at different levels

*understand format of address
and how to compute a checksum* *humans vs. whales
different wavelengths*

versus

request web page *French vs. Hungarian*

Layering

To ease software development and maximize flexibility:

- Network protocols are generally organized in **layers**
- Replace one layer without replacing surrounding layers
- Higher-level software does not have to know how to format an Ethernet packet

... or even know that Ethernet is being used

Layering

Most popular model of guiding (not specifying) protocol layers is

OSI reference model

Adopted and created by ISO

7 layers of protocols

OSI Reference Model: Layer 1

Transmits and receives raw data to communication medium

Does not care about contents voltage levels, speed, connectors

Physical

Examples: RS-232, 10BaseT

OSI Reference Model: Layer 2

Detects and corrects errors

Organizes data into packets before passing it down. Sequences packets (if necessary)

Accepts acknowledgements from receiver

Data Link

Physical

Examples: Ethernet MAC, PPP

OSI Reference Model: Layer 3

Relay and route information to destination

Manage journey of packets and figure out intermediate hops (if needed)

Network

Data Link

Physical

Examples: IP, X.25

OSI Reference Model: Layer 4

Provides a consistent interface for end-to-end (application-to-application) communication. Manages flow control

Network interface is similar to a mailbox

Transport

Network

Data Link

Physical

Examples: TCP, UDP

OSI Reference Model: Layer 5

Services to coordinate dialogue and manage data exchange

Software implemented switch

Manage multiple logical connections

Keep track of who is talking: establish & end communications

Session

Transport

Network

Data Link

Physical

Examples: HTTP 1.1, SSL, NetBIOS

OSI Reference Model: Layer 6

Presentation

Session

Transport

Network

Data Link

Physical

Data representation

Concerned with the meaning of data bits

Convert between machine representations

Examples: XDR, ASN.1, MIME

OSI Reference Model: Layer 7

Application

Presentation

Session

Transport

Network

Data Link

Physical

Collection of application-specific protocols

Examples:
email (SMTP, POP, IMAP)
file transfer (FTP)
directory services (LDAP)

Some networking terminology


Local Area Network (LAN)

Communications network

- small area (building, set of buildings)
- same, sometimes shared, transmission medium
- high data rate (often): 1 Mbps - 1 Gbps
- Low latency
- devices are peers
 - any device can initiate a data transfer with any other device

Most elements on a LAN are **workstations**

- endpoints on a LAN are called **nodes**



Connecting nodes to LANs

Adapter

- expansion slot (PCI, PC Card, USB dongle)
- usually integrated onto main board

Network adapters are referred to as **Network Interface Cards (NICs)** or **adapters** or **Network Interface Component**

Media

Wires (or RF, IR) connecting together the devices that make up a LAN

Twisted pair

- Most common:
 - STP: shielded twisted pair
 - UTP: unshielded twisted pair (e.g. Telephone cable, Ethernet 10BaseT)

Coaxial cable

- Thin (similar to TV cable)
- Thick (e.g., 10Base5, ThickNet)

Fiber

Wireless

Hubs, routers, bridges

Hub

- Device that acts as a central point for LAN cables
- Take incoming data from one port & send to all other ports

Switch

- Moves data from input to output port.
- Analyzes packet to determine destination port and makes a virtual connection between the ports.

Concentrator or repeater

- Regenerates data passing through it

Bridge

- Connects two LANs or two segments of a LAN
- Connection at data link layer (layer 2)

Router

- Determines the next network point to which a packet should be forwarded
- Connects different types of local and wide area networks at network layer (layer 3)

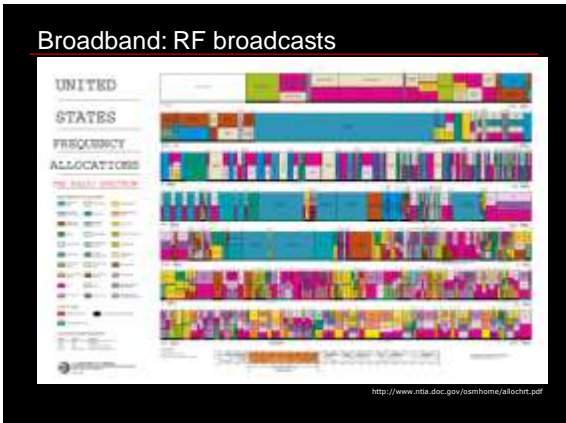
Transmission networks

Baseband

- All nodes share access to network media on an equal basis
- Data uses entire bandwidth of media

Broadband

- Data takes segment of media by dividing media into channels (frequency bands)



Broadband/Baseband: Cable TV

Broadband

55-562 MHz: analog channels 2-78
553-865 MHz: digital channels 79-136

audio

video

Baseband within Broadband

DOCSIS: Data Over Cable Service Interface Specification (approved by ITU in 1998; DOCSIS 2.0 in 2001; DOCSIS 3.0 in 2006)

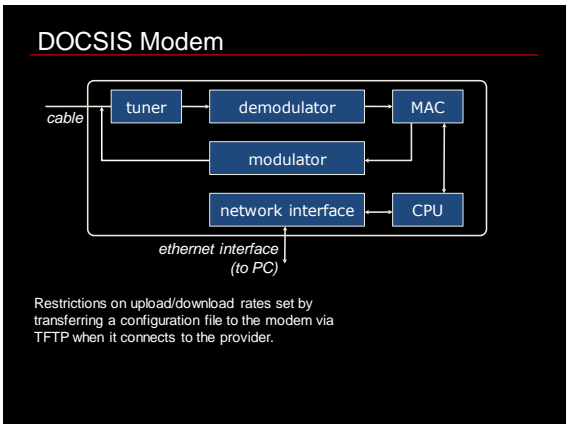
Downstream: 50-750 MHz range, 6 MHz bandwidth

- up to 38 Mbps
- received by all modems

Upstream: 5-42 MHz range

- 30.72 Mbps (10 Mbps in DOCSIS 1.0, 1.1)
- data delivered in timeslots (TDM)

DOCSIS 3.0 features **channel bonding** for greater bandwidth



Baseband: Ethernet

Standardized by IEEE as 802.3 standard

Speeds: 100 Mbps - 1 Gbps typical today

- Ethernet: 10 Mbps
- Fast Ethernet: 100 Mbps
- Gigabit Ethernet: 1 Gbps
- 10 Gbps, 100 Gbps

Network access method is

Carrier Sense Multiple Access with Collision Detection (CSMA/CD)

- Node first listens to network to see if busy
- Send
- Sense if collision occurred
- Retransmit if collision

Ethernet media

Bus topology (original design)

- originally thick coax (max 500m): 10Base5
- then... thin coax (<200m): 10Base2
 - BNC connector

Star topology (central hub or switch)

- 8 pin RJ-45 connector, UTP cable, 100 meters range
- 10BaseT for 10 Mbps
- 100BaseT for 100 Mbps
- 1000BaseT for 1 Gbps
- Cables
 - CAT-5: unshielded twisted pair
 - CAT-5e: designed for 1 Gbps
 - CAT-6: 23 gauge conductor + separator for handling crosstalk better

Wireless Ethernet media

Wireless (star topology)

- 802.11 (1-2 Mbps)
- 802.11b (11 Mbps - 4-5 Mbps realized)
- 802.11a (54 Mbps - 22-28 Mbps realized)
- 802.11g (54 Mbps - 32 Mbps realized)
- 802.11n (108 Mbps - 30-47 Mbps realized)



Client – Server Communication

Clients and Servers

- Send messages to *applications*
 - not just machines
- Client must get data to the desired *process*
 - server process must get data back to client process
- To offer a service, a server must get a **transport address** for a particular service
 - well-defined location

Machine address
versus
Transport address

Transport provider

Layer of software that accepts a network message and sends it to a remote machine

Two categories:

connection-oriented protocols

connectionless protocols

Connection-oriented Protocols

1. establish connection
2. [negotiate protocol]
3. exchange data
4. terminate connection

Connection-oriented Protocols

- | | |
|-------------------------|--------------------------------|
| 1. establish connection | <i>analogous to phone call</i> |
| 2. [negotiate protocol] | <i>dial phone number</i> |
| 3. exchange data | <i>[decide on a language]</i> |
| 4. terminate connection | <i>hang up</i> |

virtual circuit service

- provides illusion of having a dedicated circuit
- messages guaranteed to arrive in-order
- application does not have to address each message

vs. **circuit-switched service**

Connectionless Protocols

- no call setup
- send/receive data
(each packet addressed)
- no termination

Connectionless Protocols

- | | |
|--|--------------------------------|
| - no call setup | <i>analogous to mailbox</i> |
| - send/receive data
(each packet addressed) | <i>drop letter in mailbox</i> |
| - no termination | <i>(each letter addressed)</i> |

datagram service

- client is not positive whether message arrived at destination
- no state has to be maintained at client or server
- cheaper but less reliable than virtual circuit service

Ethernet

- Layers 1 & 2 of OSI model
 - Physical (1)
 - Cables: 10Base-T, 100Base-T, 1000Base-T, etc.
 - Data Link (2)
 - Ethernet bridging (via bridges)
 - Data frame parsing
 - Data frame transmission
 - Error detection
- Unreliable, connectionless communication

Ethernet

- 48-bit ethernet address
- Variable-length packet
 - 1518-byte **MTU** ← *Maximum transmission unit*
 - 18-byte header, 1500 bytes data
- Jumbo packets for Gigabit ethernet
 - 9000-byte MTU



IP – Internet Protocol

Born in 1969 as a research network of 4 machines
Funded by DoD's ARPA

Goal:

Build an efficient fault-tolerant network that could connect heterogeneous machines and link separately connected networks.

Internet Protocol

Connectionless protocol designed to handle the interconnection of a large number of local and wide-area networks that comprise the internet

IP can route from one physical network to another

IP Addressing

Each machine on an IP network is assigned a unique 32-bit number for each network interface:

- **IP address**, *not* machine address

A machine connected to several physical networks will have several IP addresses


- One for each network

IP Address space

32-bit addresses → >4 billion addresses!

- Routers would need a table of 4 billion entries
- Design routing tables so one entry can match multiple addresses
 - **hierarchy**: addresses physically close will share a common prefix

IP Addressing: networks & hosts



- first 16 bits identify Rutgers
- external routers need only one entry
 - route 128.6.*.* to Rutgers

IP Addressing: networks & hosts

- IP address
 - network #**: identifies network machine belongs to
 - host #**: identifies host on the network
- use network number to route packet to correct network
- use host number to identify specific machine

IP Addressing

Expectation:

- a few big networks and many small ones
- create different **classes** of networks
- use leading bits to identify network

class	leading bits	bits for net #	bits for host
A	0	7 (128)	24 (16M)
B	10	14 (16K)	16 (64K)
C	110	21 (2M)	8 (256)

To allow additional networks within an organization:
 use high bits of host number for a "network within a network" – **subnet**

IP Addressing

IBM: 9.0.0.0 – 9.255.255.255

00001001	xxxxxxxx	xxxxxxxx	xxxxxxxx
network #	host #		
8 bits	24 bits		

Subnet within IBM (internal routers only)

00001001	10101010	11	xxxxxxx	xxxxxxxx
network #	host #			
18 bits	14 bits			

Running out of addresses

- Huge growth
- Wasteful allocation of networks
 - Lots of unused addresses: *Does IBM need 16.7M IP addresses?*
- Every machine connected to the internet needed a worldwide-unique IP address
- Solutions: **CIDR, NAT, IPv6**

IPv6 vs. IPv4

IPv4

- 4 byte (32 bit) addresses

IPv6:

- 16-byte (128 bit) addresses
- 3.6×10^{39} possible addresses
- 8×10^{28} times more addresses than IPv4
- 4-bit priority field
- Flow label (24-bits)

IP Transport Layer Protocols

Transport-layer protocols over IP

- IP sends packets to machine
 - No mechanism for identifying sending or receiving application
- Transport layer uses a **port number** to identify the application
- TCP – Transmission Control Protocol
- UDP – User Datagram Protocol


TCP – Transmission Control Protocol

- Virtual circuit service (connection-oriented)
- Send acknowledgement for each received packet
- Checksum to validate data
- Data may be transmitted simultaneously in both directions

UDP – User Datagram Protocol

- Datagram service (connectionless)
- Data may be lost
- Data may arrive out of sequence
- Checksum for data but no retransmission
 - Bad packets dropped

IP header



vers	hlen	svc type (tos)	total length	
fragment identification		flags	fragment offset	
TTL		protocol	header checksum	
source IP address				
destination IP address				
options and pad				

Headers: TCP & UDP




src port	dest port	
seq number		
ack number		
hdr len	flags	window
checksum	urgent ptr	
options and pad		

src port	dest port
seg length	checksum

Device header (Ethernet II)



dest addr	src addr	frame type	data	CRC
-----------	----------	------------	------	-----



The End