

Computer Networking

Lent Term M/W/F 11-midday
LT1 in Gates Building

Slide Set 2

Andrew W. Moore

andrew.moore@cl.cam.ac.uk

January 2014

Topic 3: The Data Link Layer

Our goals:


- understand principles behind data link layer services:
(these are methods & mechanisms in your networking toolbox)
 - error detection, correction
 - sharing a broadcast channel: multiple access
 - link layer addressing
 - reliable data transfer, flow control: \
 - instantiation and implementation of various link layer technologies
 - Wired Ethernet (aka 802.3)
 - Wireless Ethernet (aka 802.11 WiFi)

Link Layer: Introduction

Some terminology:

- hosts and routers are **nodes**
- communication channels that connect adjacent nodes along communication path are **links**
 - wired links
 - wireless links
 - LANs
- layer-2 packet is a **frame**, encapsulates datagram

data-link layer has responsibility of transferring datagram from one node to adjacent node over a link



Link Layer (Channel) Services


- *framing, link access:*
 - encapsulate datagram into frame, adding header, trailer
 - channel access if shared medium
 - “MAC” addresses used in frame headers to identify source, dest
 - different from IP address!
- *reliable delivery between adjacent nodes*
 - we see some of this again in the Transport Topic
 - seldom used on low bit-error link (fiber, some twisted pair)
 - wireless links: high error rates
 - Q: why both link-level and end-end reliability?

Link Layer (Channel) Services - 2

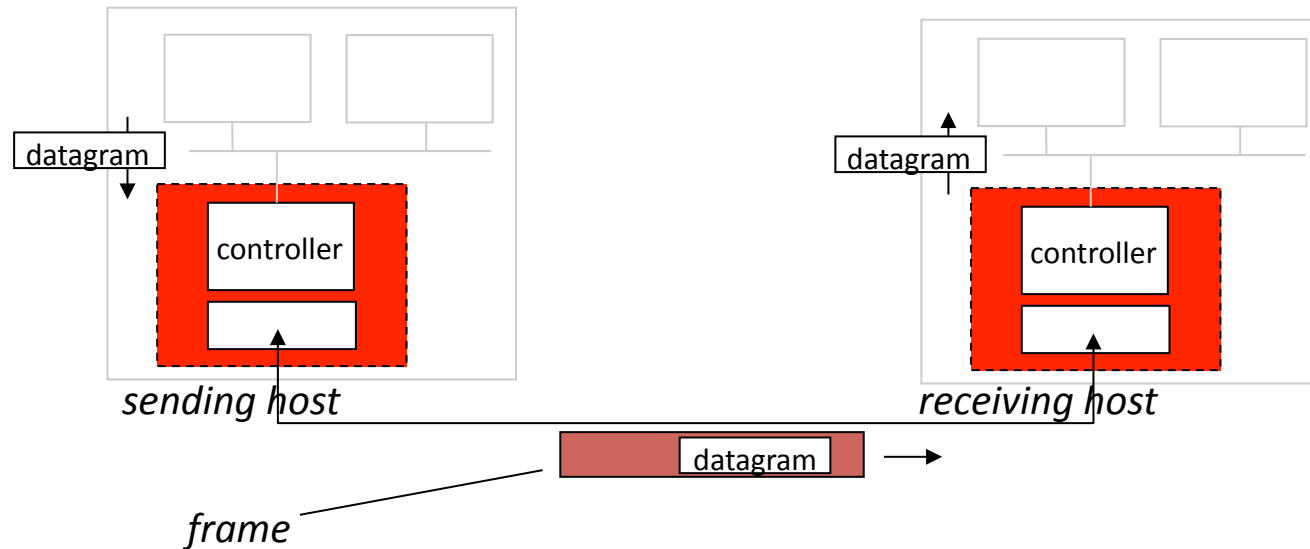
- *flow control:*
 - pacing between adjacent sending and receiving nodes
- *error detection:*
 - errors caused by signal attenuation, noise.
 - receiver detects presence of errors:
 - signals sender for retransmission or drops frame
- *error correction:*
 - receiver identifies *and corrects* bit error(s) without resorting to retransmission
- *half-duplex and full-duplex*
 - with half duplex, nodes at both ends of link can transmit, but not at same time

Where is the link layer implemented?

- in each and every host
- link layer implemented in “adaptor” (aka *network interface card* NIC)
 - Ethernet card, PCMCIA card, 802.11 card
 - implements link, physical layer
- attaches into host’s system buses
- combination of hardware, software, firmware




Adaptors Communicating



- sending side:
 - encapsulates datagram in frame
 - encodes data for the physical layer
 - adds error checking bits, provide reliability, flow control, etc.
- receiving side
 - decodes data from the physical layer
 - looks for errors, provide reliability, flow control, etc
 - extracts datagram, passes to upper layer at receiving side

Coding – a channel function

Change the representation of data.






MyPasswd



AA\$\$\$\$ff



AA\$\$\$\$ffff



MyPasswd



AA\$\$\$\$ff

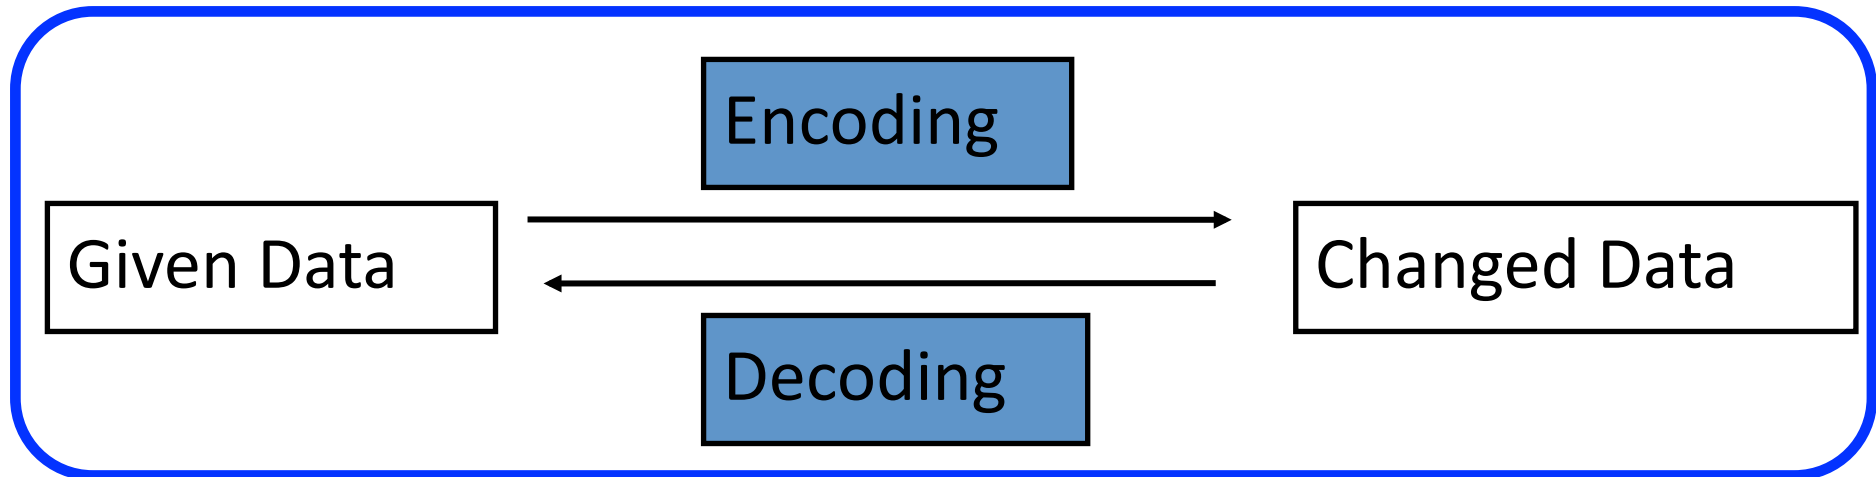



AA\$\$\$\$ffff



Coding


Change the representation of data.



1. Encryption: MyPasswd \leftrightarrow AA\$\$\$\$ff
2. Error Detection: AA\$\$\$\$ff \leftrightarrow AA\$\$\$\$ffff
3. Compression: AA\$\$\$\$ffff \leftrightarrow A2\$4f4
4. Analog: A2\$4f4 \leftrightarrow 

Line Coding Examples where Baud=bit-rate

Non-Return-to-Zero (NRZ)



Non-Return-to-Zero-Mark (NRZM) 1 = transition 0 = no transition




Non-Return-to-Zero Inverted (NRZI) (note transitions on the 1)

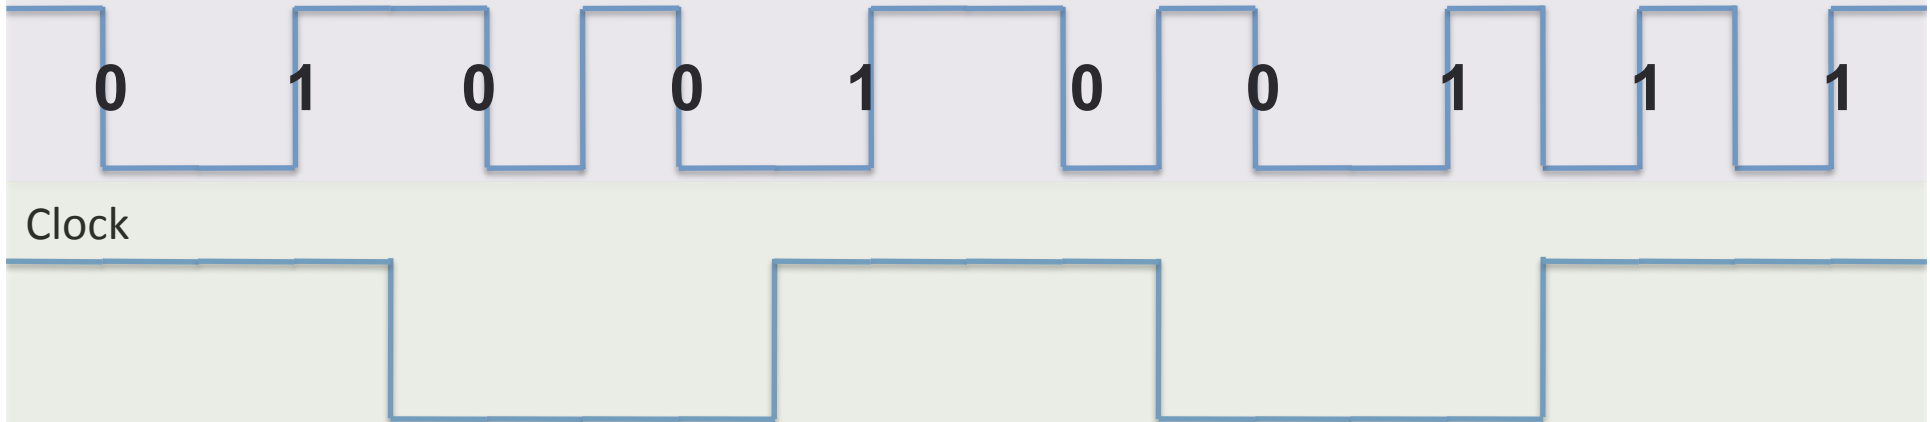


Line Coding Examples


Non-Return-to-Zero (NRZ) (Baud = bit-rate)



Manchester example (Baud = 2 x bit-rate)




Quad-level code (2 x Baud = bit-rate)



Line Coding – Block Code example

Data to send



Name	4b	5b	Description	Name	4b	5b	Description
0	0000	11110	hex data 0	Q	-NONE-	00000	Quiet
1	0001	01001	hex data 1	I	-NONE-	11111	Idle
2	0010	10100	hex data 2	J	-NONE-	11000	SSD #1
3	0011	10101	hex data 3	K	-NONE-	10001	SSD #2
4	0100	01010	hex data 4	T	-NONE-	01101	ESD #1
5	0101	01011	hex data 5	R	-NONE-	00111	ESD #2
6	0110	01110	hex data 6	H	-NONE-	00100	Halt
7	0111	01111	hex data 7				
8		1000	10010 hex data 8				
9		1001	10011 hex data 9				
A	1010	10110	hex data A				
B	1011	10111	hex data B				
C	1100	11010	hex data C				
D	1101	11011	hex data D				
E	1110	11100	hex data E				
F	1111	11101	hex data F				

Block coding transfers data with a fixed overhead: 20% less information per Baud in the case of 4B/5B

So to send data at 100Mbps; the line rate (the Baud rate) must be 125Mbps.

1Gbps uses an 8b/10b codec; encoding entire bytes at a time but with 25% overhead


Line Coding Scrambling – with secrecy

Step 1



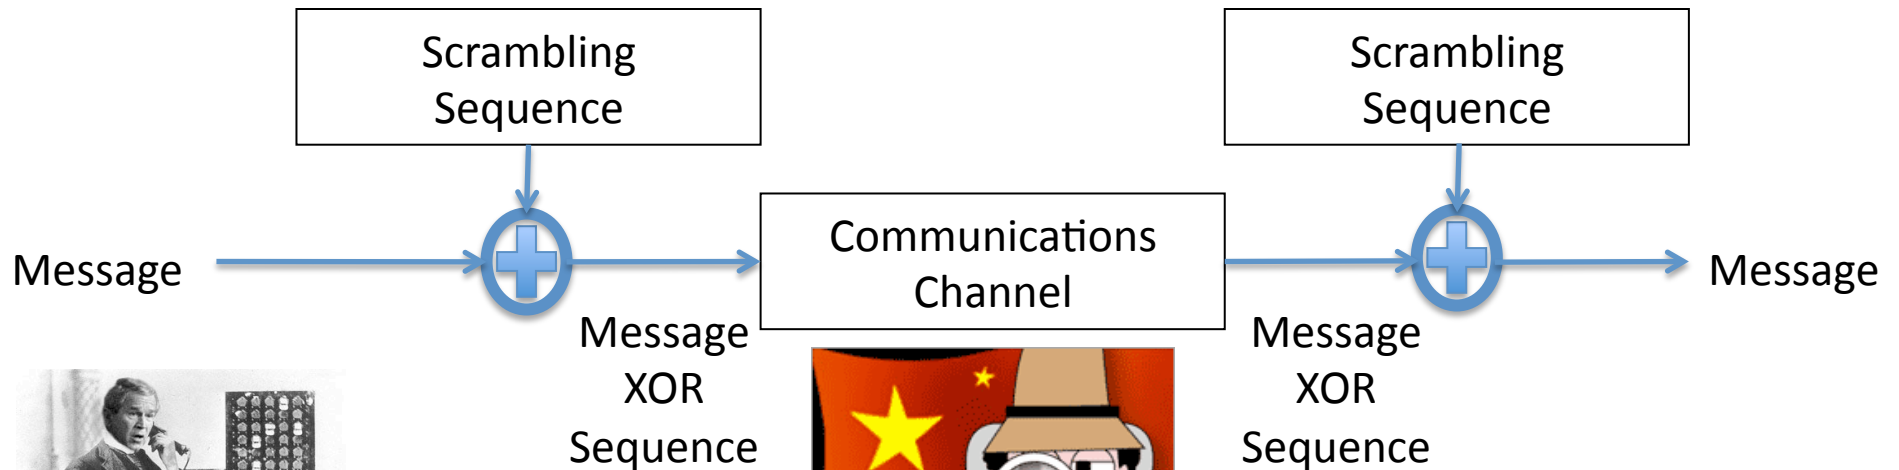


...G8wDFrB
EAFDSWbzQ7
BW2fbdTqeT
ImrukTYwQY
ndYdKb4....

**REPLICATE
SECURELY**




Step 2

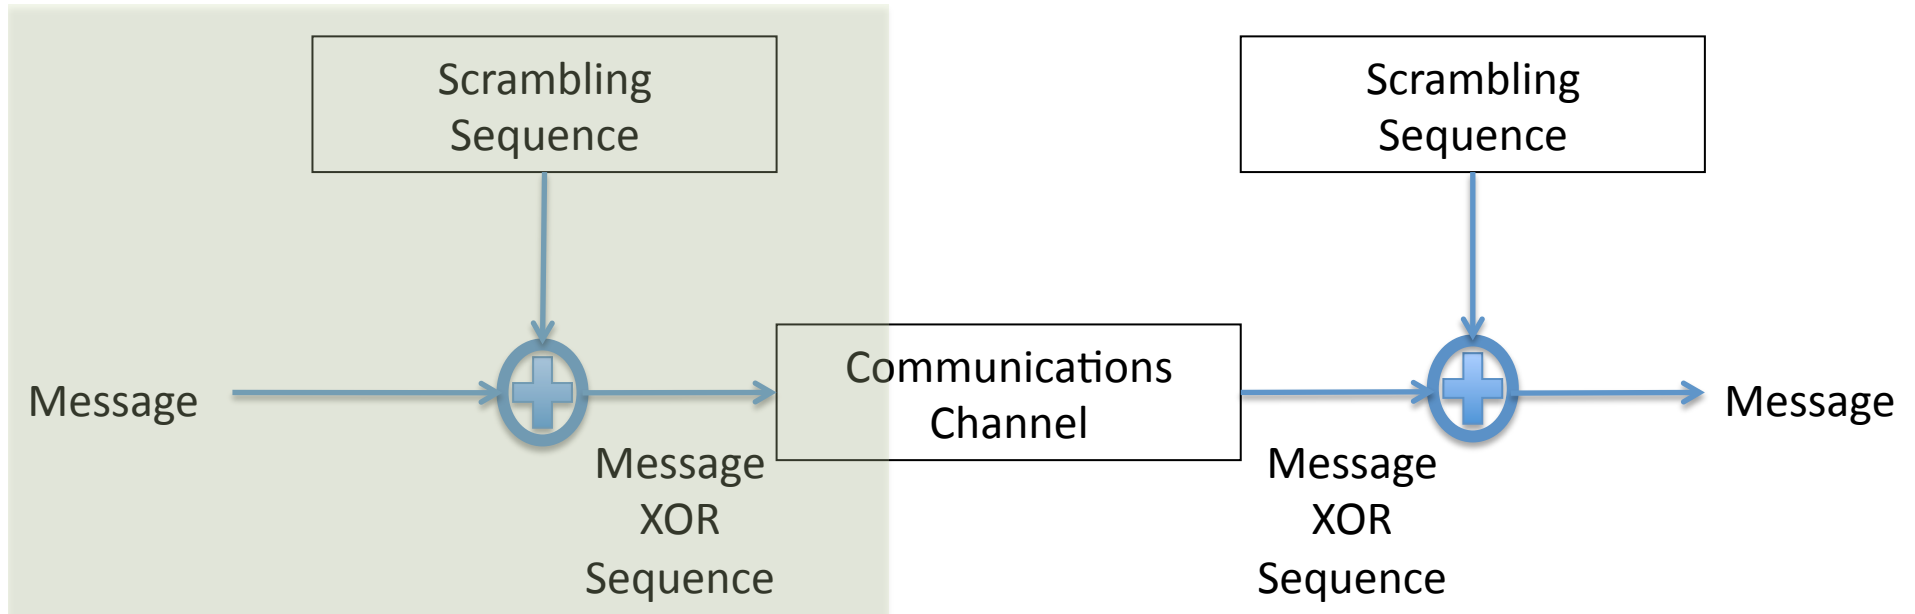


Step 3


Don't ever reuse Scrambling sequence, ever. <<< **this is quite important**



Line Coding Scrambling– no secrecy



e.g. (Self-synchronizing) scrambler




Line Coding Examples (Hybrid)

...100111101101010001000101100111010001010010110101001001110101110100...
 ...10011110110101000101000101100111010001010010110101001001110101110100...

↑
 Inserted bits marking “start of frame/block/sequence”

Scramble / Transmit / Unscramble




...0100010110011101000101001011010100100111010111010010010111011101111000...

↑
 Identify (and remove) “start of frame/block/sequence”



This gives you the Byte-delineations for *free*


64b/66b combines a scrambler and a framer. The start of frame is a pair of bits 01 or 10: 01 means “this frame is data” 10 means “this frame contains data and control” – control could be configuration information, length of encoded data or simply “this line is idle” (no data at all)



Multiple Access Mechanisms



Each dimension is orthogonal (so may be trivially combined)
There are other dimensions too; can you think of them?








Code Division Multiple Access (CDMA) (not to be confused with CSMA!)

- used in several wireless broadcast channels (cellular, satellite, etc) standards
- unique “code” assigned to each user; i.e., code set partitioning
- all users share same frequency, but each user has own “chipping” sequence (i.e., code) to encode data
- *encoded signal* = (original data) XOR (chipping sequence)
- *decoding*: inner-product of encoded signal and chipping sequence
- allows multiple users to “coexist” and transmit simultaneously with minimal interference (if codes are “orthogonal”)


CDMA Encode/Decode

sender
adds code



$$D_i = \sum_{m=1}^M Z_{i,m} \cdot c_m$$


receiver
removes code




CDMA: two-sender interference

senders

Each sender adds a **unique** code



sender removes its **unique** code



Coding Examples summary

- Common Wired coding
 - Block codecs: table-lookups
 - fixed overhead, inline control signals
 - Scramblers: shift registers
 - overhead free

Like earlier coding schemes and error correction/detection; you can combine these


- e.g, 10Gb/s Ethernet may use a hybrid

CDMA (Code Division Multiple Access)

- coping intelligently with competing sources
- Mobile phones

Error Detection and Correction

How to use coding to deal with errors in data communication?




Basic Idea :

1. Add additional information to a message.
2. Detect an error and re-send a message.
Or, fix an error in the received message.

Error Detection and Correction

How to use coding to deal with errors in data communication?



Basic Idea :


1. Add additional information to a message.
2. Detect an error and re-send a message.
Or, fix an error in the received message.

Error Detection

EDC= Error Detection and Correction bits (redundancy = overhead)

D = Data protected by error checking, may include header fields

- Error detection not 100% reliable!
 - protocol may miss some errors, but rarely
 - larger EDC field yields better detection and correction





Error Detection Code

Sender:

```
Y = generateCheckBit(X);  
send(XY);
```


Receiver:

```
receive(X1Y1);  
Y2=generateCheckBit(X1);  
if (Y1 != Y2) ERROR;  
else NOERROR
```



Error Detection Code: Parity

Add one bit, such that the number of 1's is even.



0000	0
------	---

0001	1
------	---

1001	0
------	---

✗

0001	0
------	---

✓

0001	1
------	---

✓


1111	0
------	---

Problem: This simple parity cannot detect two-bit errors.

Parity Checking


Single Bit Parity:

Detect single bit errors



Two Dimensional Bit Parity:

Detect *and correct* single bit errors



1	0	1	0	1	1
1	1	1	1	0	0
0	1	1	1	0	1
0	0	1	0	1	0

no errors

1	0	1	0	1	1
1	0	1	1	0	0
0	1	1	1	0	1
0	0	1	0	1	0

parity error

*correctable
single bit error*

Internet checksum

Goal: detect “errors” (e.g., flipped bits) in transmitted packet
(note: used at transport layer only)

Sender:

- treat segment contents as sequence of 1bit integers
- checksum: addition (1's complement sum) of segment contents
- sender puts checksum value into UDP checksum field


Receiver:

- compute checksum of received segment
- check if computed checksum equals checksum field value:
 - NO - error detected
 - YES - no error detected. *But maybe errors nonetheless?*

Error Detection Code: CRC

- CRC means “Cyclic Redundancy Check”.
- More powerful than parity.
 - It can detect various kinds of errors, including 2-bit errors.
- More complex: multiplication, binary division.
- Parameterized by n -bit divisor P .
 - Example: 3-bit divisor 101.
 - Choosing good P is crucial.

CRC with 3-bit Divisor 101



111
100

same check bits from Parity,
but different ones from CRC

Multiplication by 2^3
 $D2 = D * 2^3$

Binary Division by 101
 $CheckBit = (D2) \text{ rem } (101)$

Add three 0's at the end

Kurose p478 §5.2.3
Peterson p97 §2.4.3³³

The divisor (**G**) – Secret sauce of CRC


- If the divisor were 100, instead of 101, data 1111 and 1001 would give the same check bit 00.
- Mathematical analysis about the divisor:
 - Last bit should be 1.
 - Should contain at least two 1's.
 - Should be divisible by 11.
- ATM, HDLC, Ethernet each use a CRC with well-chosen fixed divisors

Divisor analysis keeps mathematicians in jobs
(a branch of *pure* math: combinatorial mathematics)

Checksumming: Cyclic Redundancy Check recap

- view data bits, **D**, as a binary number
- choose $r+1$ bit pattern (generator), **G**
- goal: choose r CRC bits, **R**, such that
 - $\langle D, R \rangle$ exactly divisible by G (modulo 2)
 - receiver knows G , divides $\langle D, R \rangle$ by G . If non-zero remainder: error detected!
 - can detect all burst errors less than $r+1$ bits
- widely used in practice (Ethernet, 802.11 WiFi, ATM)

← d bits → ← r bits →



$D * 2^r \text{ XOR } R$ *mathematical formula*

CRC Another Example – this time with long division

Want:

$$D \cdot 2^r \text{ XOR } R = nP$$


equivalently:

$$D \cdot 2^r = nP \text{ XOR } R$$

equivalently:

if we divide $D \cdot 2^r$ by P ,
want remainder R

$$R = \text{remainder} \left[\frac{D \cdot 2^r}{P} \right]$$





Error Detection Code becomes....

Sender:

```
Y = generateCheckBit(X);  
send(XY);
```

Receiver:

```
receive(X1Y1);  
Y2=generateCheckBit(X1);  
if (Y1 != Y2) ERROR;  
else NOERROR
```






Forward Error Correction (FEC)

Sender:

```
Y = generateCheckBit(X);  
send(XY);
```

Receiver:

```
receive(X1Y1);  
Y2=generateCheckBit(X1);  
if (Y1 != Y2) FIXERROR(X1Y1);  
else NOERROR
```





Forward Error Correction (FEC)

Sender:

```
Y = generateCheckBit(X);  
send(XY);
```

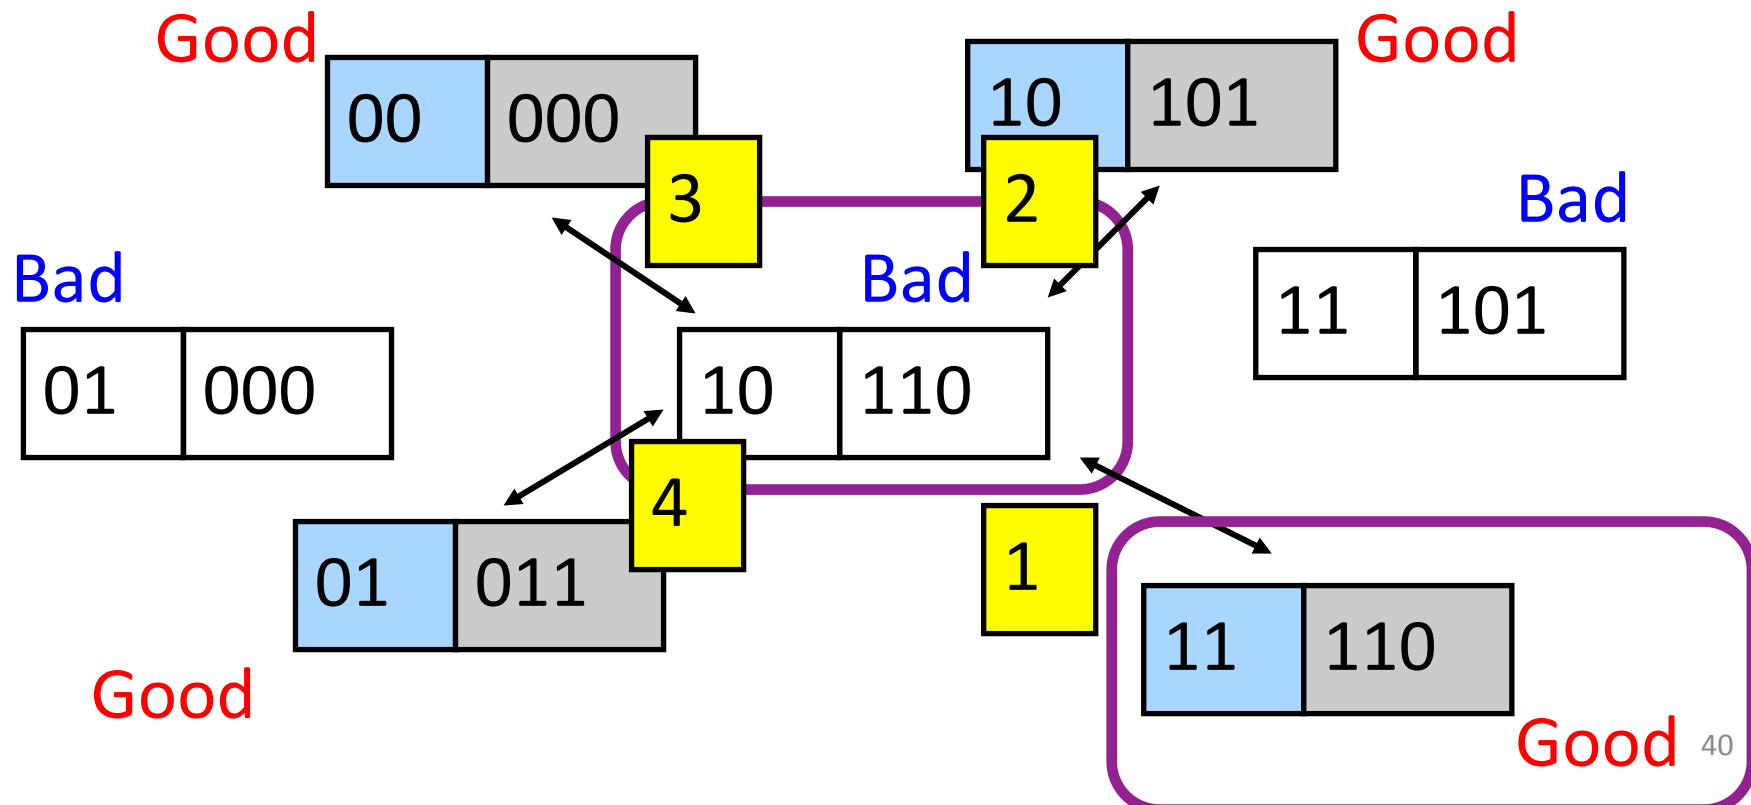
Receiver:

```
receive(X1Y1);  
Y2=generateCheckBit(X1);  
if (Y1 != Y2) FIXERROR(X1Y1);  
else NOERROR
```



Basic Idea of Forward Error Correction

Replace erroneous data by its “closest” error-free data.



Error Detection vs Correction

Error Correction:

- Cons: More check bits. False recovery.
- Pros: No need to re-send.

Error Detection:

- Cons: Need to re-send.
- Pros: Less check bits.


Usage:

- Correction: A lot of noise. Expensive to re-send.
- Detection: Less noise. Easy to re-send.
- Can be used together.

Multiple Access Links and Protocols

Two types of “links”:


- point-to-point
 - point-to-point link between Ethernet switch and host
- **broadcast** (shared wire or medium)
 - old-fashioned wired Ethernet (*here be dinosaurs* – extinct)
 - upstream HFC (Hybrid Fiber-Coax – the Coax may be broadcast)
 - Home plug / Powerline networking
 - 802.11 wireless LAN



shared wire (e.g.,
cabled Ethernet)



shared RF
(e.g., 802.11 WiFi)



shared RF
(satellite)



humans at a
cocktail party
(shared air, acoustical)

Multiple Access protocols

- single shared broadcast channel
 - two or more simultaneous transmissions by nodes:
interference
 - **collision** if node receives two or more signals at the same time
- multiple access protocol*
- distributed algorithm that determines how nodes share channel, i.e., determine when node can transmit
 - communication about channel sharing must use channel itself!
 - no out-of-band channel for coordination

Ideal Multiple Access Protocol

Broadcast channel of rate R bps

1. when one node wants to transmit, it can send at rate R
2. when M nodes want to transmit, each can send at average rate R/M
3. fully decentralized:
 - no special node to coordinate transmissions
 - no synchronization of clocks, slots
4. simple

MAC Protocols: a taxonomy

Three broad classes:


- **Channel Partitioning**
 - divide channel into smaller “pieces” (time slots, frequency, code)
 - allocate piece to node for exclusive use
- **Random Access**
 - channel not divided, allow collisions
 - “recover” from collisions
- **“Taking turns”**
 - nodes take turns, but nodes with more to send can take longer turns



Channel Partitioning MAC protocols: TDMA *(time travel warning – we mentioned this earlier)*

TDMA: time division multiple access

- access to channel in "rounds"
- each station gets fixed length slot (length = pkt trans time) in each round
- unused slots go idle
- example: station LAN, 1,3,4 have pkt, slots 2,5,6 idle






Channel Partitioning MAC protocols: FDMA *(time travel warning – we mentioned this earlier)*

FDMA: frequency division multiple access

- channel spectrum divided into frequency bands
- each station assigned fixed frequency band
- unused transmission time in frequency bands go idle
- example: station LAN, 1,3,4 have pkt, frequency bands 2,5,6 idle



“Taking Turns” MAC protocols

channel partitioning MAC protocols:

- share channel *efficiently* and *fairly* at high load
- inefficient at low load: delay in channel access, $1/N$ bandwidth allocated even if only 1 active node!

Random access MAC protocols

- efficient at low load: single node can fully utilize channel
- high load: collision overhead


“taking turns” protocols

look for best of both worlds!

“Taking Turns” MAC protocols

Polling:


- master node “invites” slave nodes to transmit in turn
- typically used with “dumb” slave devices
- concerns:
 - polling overhead
 - latency
 - single point of failure (master)



“Taking Turns” MAC protocols

Token passing:


- ❑ control **token** passed from one node to next sequentially.
- ❑ token message
- ❑ concerns:
 - token overhead
 - latency
 - single point of failure (token)
- concerns fixed in part by a slotted ring (many simultaneous *tokens*)



Cambridge students – this is YOUR heritage
Cambridge RING, Cambridge Fast RING,
Cambridge Backbone RING, these things gave us ATDM (and ATM)

ATM

In TDM a sender may only use a pre-allocated slot



In ATM a sender transmits labeled cells whenever necessary



ATM = Asynchronous Transfer Mode – an ugly expression
think of it as ATDM – Asynchronous Time Division Multiplexing

That's a variant of **PACKET SWITCHING** to the rest of us – just like Ethernet
but using fixed length slots/packets/cells

Use the media when you need it, but

ATM had virtual circuits and these needed setup....

Worse ATM had an utterly irrational size

Random Access MAC Protocols

- When node has packet to send
 - Transmit at full channel data rate
 - No *a priori* coordination among nodes
- Two or more transmitting nodes \Rightarrow collision
 - Data lost
- Random access MAC protocol specifies:
 - How to detect collisions
 - How to recover from collisions
- Examples
 - ALOHA and Slotted ALOHA
 - CSMA, CSMA/CD, CSMA/CA (wireless)

Key Ideas of Random Access

- **Carrier sense**
 - *Listen before speaking, and don't interrupt*
 - Checking if someone else is already sending data
 - ... and waiting till the other node is done
- **Collision detection**
 - *If someone else starts talking at the same time, stop*
 - Realizing when two nodes are transmitting at once
 - ...by detecting that the data on the wire is garbled
- **Randomness**
 - *Don't start talking again right away*
 - Waiting for a random time before trying again

CSMA (Carrier Sense Multiple Access)

- CSMA: **listen** before transmit
 - If channel sensed idle: transmit entire frame
 - If channel sensed busy, defer transmission
- Human analogy: don't interrupt others!
- Does this eliminate all collisions?
 - No, because of nonzero propagation delay

CSMA Collisions


Propagation delay: two nodes may not hear each other's before sending.

Would slots hurt or help?

CSMA reduces but does not eliminate collisions

Biggest remaining problem?

Collisions still take full slot!
How do you fix that?




CSMA/CD (Collision Detection)

- CSMA/CD: carrier sensing, deferral as in CSMA
 - **Collisions detected within short time**
 - Colliding transmissions aborted, reducing wastage
- Collision detection easy in wired LANs:
 - Compare transmitted, received signals
- Collision detection difficult in wireless LANs:
 - Reception shut off while transmitting (well, perhaps not)
 - Not perfect broadcast (limited range) so collisions local
 - Leads to use of *collision avoidance* instead (later)

CSMA/CD Collision Detection

B and D can tell that collision occurred.

Note: for this to work, need restrictions on minimum frame size and maximum distance. Why?



Limits on CSMA/CD Network Length



- Latency depends on physical length of link
 - Time to propagate a packet from one end to the other
- Suppose *A* sends a packet at time t
 - And *B* sees an idle line at a time just before $t+d$
 - ... so *B* happily starts transmitting a packet
- *B* detects a collision, and sends **jamming signal**
 - But *A* can't see collision until $t+2d$

Performance of CSMA/CD

- Time wasted in collisions
 - Proportional to distance d
- Time spend transmitting a packet
 - Packet length p divided by bandwidth b
- Rough estimate for efficiency (K some constant)

$$E \sim \frac{\frac{p}{b}}{\frac{p}{b} + Kd}$$

- Note:
 - For large packets, small distances, $E \sim 1$
 - As bandwidth increases, E decreases
 - That is why high-speed LANs are all switched


Benefits of Ethernet

- Easy to administer and maintain
- Inexpensive
- Increasingly higher speed
- Evolvable!

Evolution of Ethernet


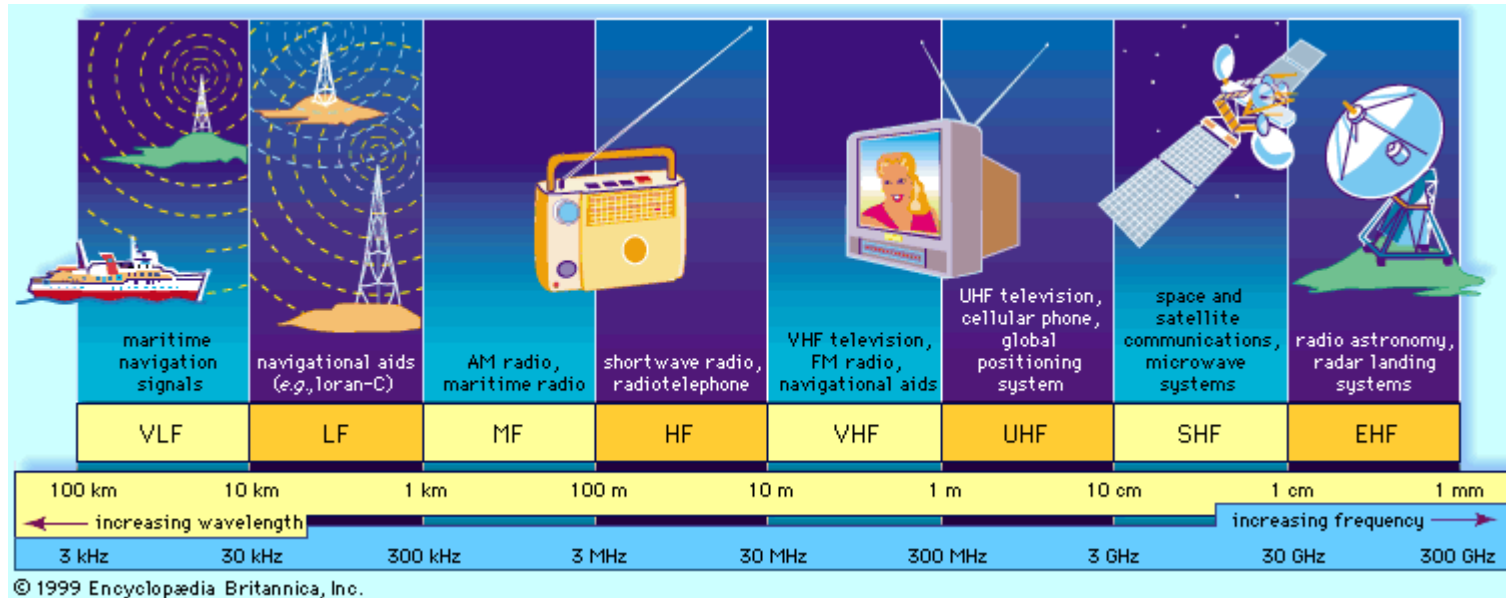
- Changed **everything** except the frame **format**
 - From single coaxial cable to hub-based star
 - From shared media to **switches**
 - From electrical signaling to optical
- **Lesson #1**
 - The right **interface** can accommodate many **changes**
 - Implementation is hidden behind interface
- **Lesson #2**
 - Really hard to displace the dominant technology
 - Slight performance improvements are not enough

Ethernet: CSMA/CD Protocol



- **Carrier sense:** wait for link to be idle
- **Collision detection:** listen while transmitting
 - No collision: transmission is complete
 - Collision: abort transmission & send **jam** signal
- **Random access:** **binary exponential back-off**
 - After collision, wait a random time before trying again
 - After m^{th} collision, choose K randomly from $\{0, \dots, 2^m - 1\}$
 - ... and wait for $K * 512$ bit times before trying again
 - Using min packet size as “slot”
 - **If transmission occurring when ready to send, wait until end of transmission (CSMA)**

The Wireless Spectrum



Metrics for evaluation / comparison of wireless technologies

- Bitrate or Bandwidth
- Range - PAN, LAN, MAN, WAN
- Two-way / One-way
- Multi-Access / Point-to-Point
- Digital / Analog
- Applications and industries
- Frequency – Affects most physical properties:
 - Distance (free-space loss)
 - Penetration, Reflection, Absorption
 - Energy proportionality
 - Policy: Licensed / Deregulated
 - Line of Sight (Fresnel zone)
 - Size of antenna
- Determined by wavelength – $\lambda = \frac{v}{f}$,

Wireless Communication Standards

- Cellular (**800/900/1700/1800/1900Mhz**):
 - 2G: GSM / CDMA / GPRS /EDGE
 - 3G: CDMA2000/UMTS/HSDPA/EVDO
 - 4G: LTE, WiMax
- IEEE 802.11 (aka WiFi):
 - b: **2.4Ghz** band, 11Mbps (*~4.5 Mbps operating rate*)
 - g: **2.4Ghz**, 54-108Mbps (*~19 Mbps operating rate*)
 - a: **5.0Ghz** band, 54-108Mbps (*~25 Mbps operating rate*)
 - n: **2.4/5Ghz**, 150-600Mbps (4x4 mimo).
- IEEE 802.15 – lower power wireless:
 - 802.15.1: **2.4Ghz**, 2.1 Mbps (Bluetooth)
 - 802.15.4: **2.4Ghz**, 250 Kbps (Sensor Networks)


What Makes Wireless Different?

- Broadcast and multi-access medium...
 - err, so....
- BUT, Signals sent by sender don't always end up at receiver intact
 - Complicated physics involved, which we won't discuss
 - But what can go wrong?

Path Loss / Path Attenuation

- Free Space Path Loss:
$$\text{FSPL} = \left(\frac{4\pi d}{\lambda} \right)^2$$
$$= \left(\frac{4\pi d f}{c} \right)^2$$
 - d = distance
 - λ = wave length
 - f = frequency
 - c = speed of light
- Reflection, Diffraction, Absorption
- Terrain contours (Urban, Rural, Vegetation).
- Humidity

Multipath Effects



- Signals bounce off surface and interfere with one another
- Self-interference

Interference from Other Sources

- External Interference
 - Microwave is turned on and blocks your signal
 - Would that affect the sender or the receiver?
- Internal Interference
 - Hosts within range of each other collide with one another's transmission
- We have to tolerate path loss, multipath, etc., but we can try to avoid internal interference

Wireless Bit Errors

- The lower the SNR (Signal/Noise) the higher the Bit Error Rate (BER)
- We could make the signal stronger...
- Why is this not always a good idea?
 - Increased signal strength requires more power
 - Increases the interference range of the sender, so you interfere with more nodes around you
 - And then they increase their power.....
- Local link-layer Error Correction schemes can correct **some** problems

Lets focus on 802.11

aka - WiFi ...


What makes it special?

Deregulation > Innovation > Adoption > Lower cost = Ubiquitous technology

JUST LIKE ETHERNET – not lovely but sufficient

802.11 Architecture

802.11 frames exchanges



802.3 (Ethernet)
frames exchanged


Figure 6.7 ♦ IEEE 802.11 LAN architecture

- Designed for limited area
- AP' s (Access Points) set to specific channel
- Broadcast beacon messages with SSID (Service Set Identifier) and MAC Address periodically
- Hosts scan all the channels to discover the AP' s
 - Host associates with AP

Wireless Multiple Access Technique?


- Carrier Sense?
 - Sender can listen before sending
 - What does that tell the sender?
- Collision Detection?
 - Where do collisions occur?
 - How can you detect them?

Hidden Terminals



- A and C can both send to B but **can't hear each other**
 - A is a *hidden terminal* for C and vice versa
- Carrier Sense will be **ineffective**

Exposed Terminals



- **Exposed node:** B sends a packet to A; C hears this and decides not to send a packet to D (despite the fact that this will not cause interference)!
- Carrier sense would prevent a successful transmission.


Key Points

- No concept of a global collision
 - Different receivers hear different signals
 - Different senders reach different receivers
- Collisions are at receiver, not sender
 - Only care if receiver can hear the sender clearly
 - It does not matter if sender can hear someone else
 - As long as that signal does not interfere with receiver
- Goal of protocol:
 - Detect if receiver can hear sender
 - Tell senders who might interfere with receiver to shut up

Basic Collision Avoidance


- Since can't detect collisions, we try to *avoid* them
- Carrier sense:
 - When medium busy, choose random interval
 - Wait that many **idle** timeslots to pass before sending
- When a collision is inferred, retransmit with binary exponential backoff (like Ethernet)
 - Use **ACK** from receiver to infer “no collision”
 - Use exponential backoff to adapt contention window

CSMA/CA -MA with Collision Avoidance




- Before every data transmission
 - Sender sends a Request to Send (RTS) frame containing the length of the transmission
 - Receiver respond with a Clear to Send (CTS) frame
 - Sender sends data
 - Receiver sends an ACK; now another sender can send data
- When sender doesn't get a CTS back, it assumes collision

CSMA/CA, con' t



- If other nodes hear RTS, but not CTS: **send**
 - Presumably, destination for first sender is out of node's range ...


CSMA/CA, con' t



- If other nodes hear RTS, but not CTS: **send**
 - Presumably, destination for first sender is out of node's range ...
 - ... Can cause problems when a CTS is **lost**
- When you hear a CTS, you keep quiet until scheduled transmission is over (hear ACK)

RTS / CTS Protocols (CSMA/CA)

B sends to C




Overcome hidden terminal problems with contention-free protocol

1. B sends to C **Request To Send** (RTS)
2. A hears RTS and defers (to allow C to answer)
3. C replies to B with **Clear To Send** (CTS)
4. D hears CTS and defers to allow the data
5. B sends to C


Preventing Collisions Altogether

- Frequency Spectrum partitioned into several channels
 - Nodes within interference range can use separate channels



- Now A and C can send without any interference!
- Most cards have only 1 transceiver
 - **Not Full Duplex: Cannot send and receive at the same time**
 - Aggregate Network throughput doubles

CSMA/CA and RTS/CTS



RTS/CTS

- helps with hidden terminal
- good for high-traffic Access Points
- often turned on/off dynamically

Without RTS/CTS

- lower latency -> faster!
- reduces wasted b/w
 - if the $Pr(\text{collision})$ is low
- good for when net is small and not *weird*
 - eg no hidden/exposed terminals

CSMA/CD vs CSMA/CA (without RTS/CTS)

CD Collision Detect

wired – listen and talk

1. Listen for others
2. Busy? goto 1.
3. Send message (and listen)
4. Collision?
 - a. JAM
 - b. increase your BEB
 - c. sleep
 - d. goto 1.

CA Collision Avoidance


wireless – talk OR listen

1. Listen for others
2. Busy?
 - a. increase your BEB
 - b. sleep
 - c. goto 1.
3. Send message
4. Wait for ACK (*MAC ACK*)
5. Got No ACK from MAC?
 - a. increase your BEB
 - b. sleep
 - c. goto 1.

Changing the rules: an 802.11 feature

Rate Adaptation

- base station, mobile dynamically change transmission rate (physical layer modulation technique) as mobile moves, SNR varies



- QAM256 (8 Mbps)
- - - QAM16 (4 Mbps)
- BPSK (1 Mbps)
- operating point

1. SNR decreases, BER increase as node moves away from base station

2. When BER becomes too high, switch to lower transmission rate but with lower BER

Summary of MAC protocols

- *channel partitioning*, by time, frequency or code
 - Time Division, Frequency Division
- *random access* (dynamic),
 - ALOHA, S-ALOHA, CSMA, CSMA/CD
 - carrier sensing: easy in some technologies (wire), hard in others (wireless)
 - CSMA/CD used in Ethernet
 - CSMA/CA used in 802.11
- *taking turns*
 - polling from central site, token passing
 - Bluetooth, FDDI, IBM Token Ring

MAC Addresses

- MAC (or LAN or physical or Ethernet) address:
 - function: *get frame from one interface to another physically-connected interface (same network)*
 - 48 bit MAC address (for most LANs)
 - *burned* in NIC ROM, nowadays usually software settable and set at boot time

```
awm22@rio:~$ ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:30:48:fe:c0:64
          inet addr:128.232.33.4  Bcast:128.232.47.255  Mask:255.255.240.0
          inet6 addr: fe80::230:48ff:fefe:c064/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:215084512 errors:252 dropped:25 overruns:0 frame:123
          TX packets:146711866 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:170815941033 (170.8 GB)  TX bytes:86755864270 (86.7 GB)
          Memory:f0000000-f0020000
```

LAN Address (more)


- MAC address allocation administered by IEEE
- manufacturer buys portion of MAC address space (to assure uniqueness)
- analogy:
 - (a) MAC address: like Social Security Number
 - (b) IP address: like postal address
- MAC flat address → portability
 - can move LAN card from one LAN to another
- IP hierarchical address NOT portable
 - address depends on IP subnet to which node is attached


Hubs

... physical-layer (“dumb”) repeaters:

- bits coming in one link go out *all* other links at same rate
- all nodes connected to hub can collide with one another
- no frame buffering
- no CSMA/CD at hub: host NICs detect collisions

Collision Domain
in CSMA/CD *speaks*






CSMA/CD Lives....

BRAINS... BRAINS...



Home Plug and similar Powerline Networking....




Switch

(like a Hub but smarter)

- **link-layer device: smarter than hubs, take *active* role**
 - store, forward Ethernet frames
 - examine incoming frame's MAC address, **selectively** forward frame to one-or-more outgoing links when frame is to be forwarded on segment, uses CSMA/CD to access segment
- ***transparent***
 - hosts are unaware of presence of switches
- ***plug-and-play, self-learning***
 - switches do not need to be configured

Switch: allows *multiple* simultaneous transmissions


- hosts have dedicated, direct connection to switch
- switches buffer packets
- Ethernet protocol used on *each* incoming link, but no collisions; full duplex
 - each link is its own collision domain
- **switching:** A-to-A' and B-to-B' simultaneously, without collisions
 - not possible with dumb hub



switch with six interfaces
(1,2,3,4,5,6)

Switch Table


- Q: how does switch know that A' reachable via interface 4, B' reachable via interface 5?
- A: each switch has a **switch table**, each entry:
 - (MAC address of host, interface to reach host, time stamp)
- looks like a routing table!
- Q: how are entries created, maintained in switch table?
 - something like a routing protocol?



*switch with six interfaces
(1,2,3,4,5,6)*

Switch: self-learning (recap)

- switch *learns* which hosts can be reached through which interfaces
 - when frame received, switch “learns” location of sender: incoming LAN segment
 - records sender/location pair in switch table



MAC addr	interface	TTL
A	1	60

Switch table
(initially empty)

Switch: frame filtering/forwarding

When frame received:


1. record link associated with sending host
2. index switch table using MAC dest address
- 3. if** entry found for destination
 then {
 if dest on segment from which frame arrived
 then drop the frame
 else forward the frame on interface indicated
 }
 else flood

*forward on all but the interface
on which the frame arrived*



Self-learning, forwarding: example

- frame destination unknown: *flood*
- destination A location known: *selective send*




MAC addr	interface	TTL
A	1	60
A'	4	60

Switch table
(initially empty)

Interconnecting switches


- switches can be connected together



- ❑ **Q:** sending from A to G - how does S₁ know to forward frame destined to F via S₄ and S₃?
- ❑ **A:** self learning! (works exactly the same as in single-switch case – **flood/forward/drop**)

Flooding Can Lead to Loops


- Flooding can lead to **forwarding loops**
 - E.g., if the network contains a cycle of switches
 - “Broadcast storm”





Solution: Spanning Trees

- Ensure the forwarding **topology** has no loops
 - Avoid using some of the links when flooding
 - ... to prevent loop from forming
- **Spanning tree**
 - **Sub-graph** that covers all vertices but *contains no cycles*
 - Links not in the spanning tree do not forward frames

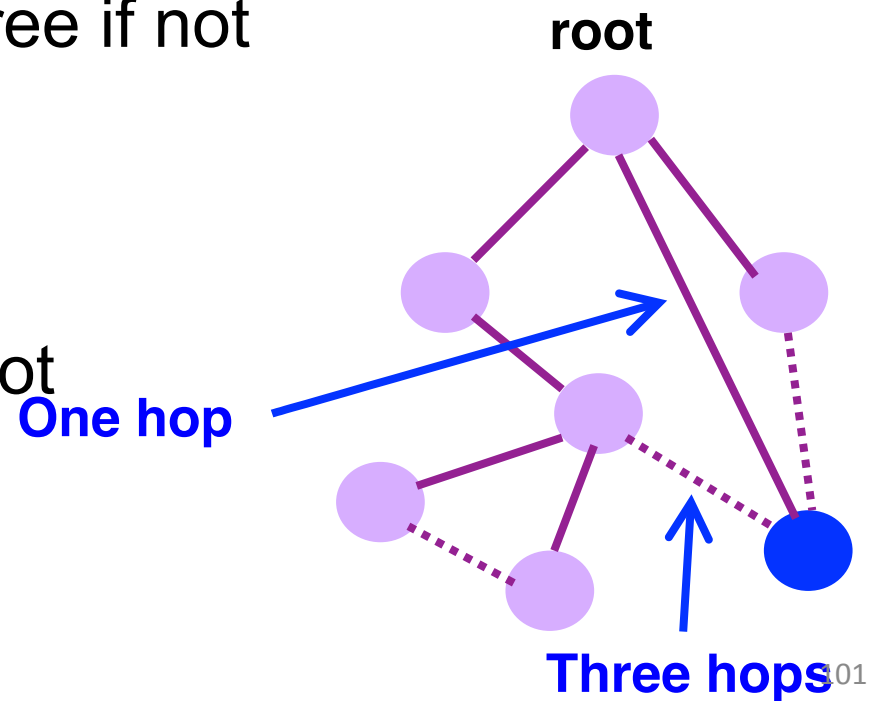


What Do We Know?

- Shortest paths to (or from) a node form a tree
- So, algorithm has two aspects :
 - Pick a root
 - Compute shortest paths to it
- Only keep the links on shortest-path

Constructing a Spanning Tree

- Switches need to **elect** a **root**
 - The switch w/ smallest identifier (MAC addr)
- Each switch determines if each interface is on the **shortest path** from the root
 - Excludes it from the tree if not
- Messages (Y, d, X)
 - From node X
 - Proposing Y as the root
 - And the distance is d




Steps in Spanning Tree Algorithm

- Initially, each switch proposes itself as the root
 - Switch sends a message out every interface
 - ... proposing itself as the root with distance 0
 - Example: switch X announces (X, 0, X)
- Switches update their view of the root
 - Upon receiving message (Y, d, Z) from Z, check Y's id
 - If new id smaller, start viewing that switch as root
- Switches compute their distance from the root
 - Add 1 to the distance received from a neighbor
 - Identify interfaces not on shortest path to the root
 - ... and exclude them from the spanning tree
- If root or shortest distance to it **changed**, “flood” updated message (Y, d+1, X)


Example From Switch #4' s Viewpoint

- Switch #4 thinks it is the root
 - Sends (4, 0, 4) message to 2 and 7
- Then, switch #4 hears from #2
 - Receives (2, 0, 2) message from 2
 - ... and thinks that #2 is the root
 - And realizes it is just one hop away
- Then, switch #4 hears from #7
 - Receives (2, 1, 7) from 7
 - And realizes this is a longer path
 - So, prefers its own one-hop path
 - And removes 4-7 link from the tree



Example From Switch #4' s Viewpoint

- Switch #2 hears about switch #1
 - Switch 2 hears (1, 1, 3) from 3
 - Switch 2 starts treating 1 as root
 - And sends (1, 2, 2) to neighbors
- Switch #4 hears from switch #2
 - Switch 4 starts treating 1 as root
 - And sends (1, 3, 4) to neighbors
- Switch #4 hears from switch #7
 - Switch 4 receives (1, 3, 7) from 7
 - And realizes this is a longer path
 - So, prefers its own three-hop path
 - And removes 4-7 link from the tree



Robust Spanning Tree Algorithm

- Algorithm must react to **failures**
 - Failure of the root node
 - Need to elect a new root, with the next lowest identifier
 - Failure of other switches and links
 - Need to recompute the spanning tree
- Root switch continues sending messages
 - Periodically reannouncing itself as the root (1, 0, 1)
 - Other switches continue forwarding messages
- Detecting failures through timeout (**soft state**)
 - If no word from root, times out and claims to be the root
 - Delay in reestablishing spanning tree is **major problem**
 - Work on rapid spanning tree algorithms...

Topic 3: Summary

- principles behind data link layer services:
 - error detection, correction
 - sharing a broadcast channel: multiple access
 - link layer addressing
- instantiation and implementation of various link layer technologies
 - Ethernet
 - switched LANs
 - WiFi
- algorithms
 - Binary Exponential Backoff
 - Spanning Tree