
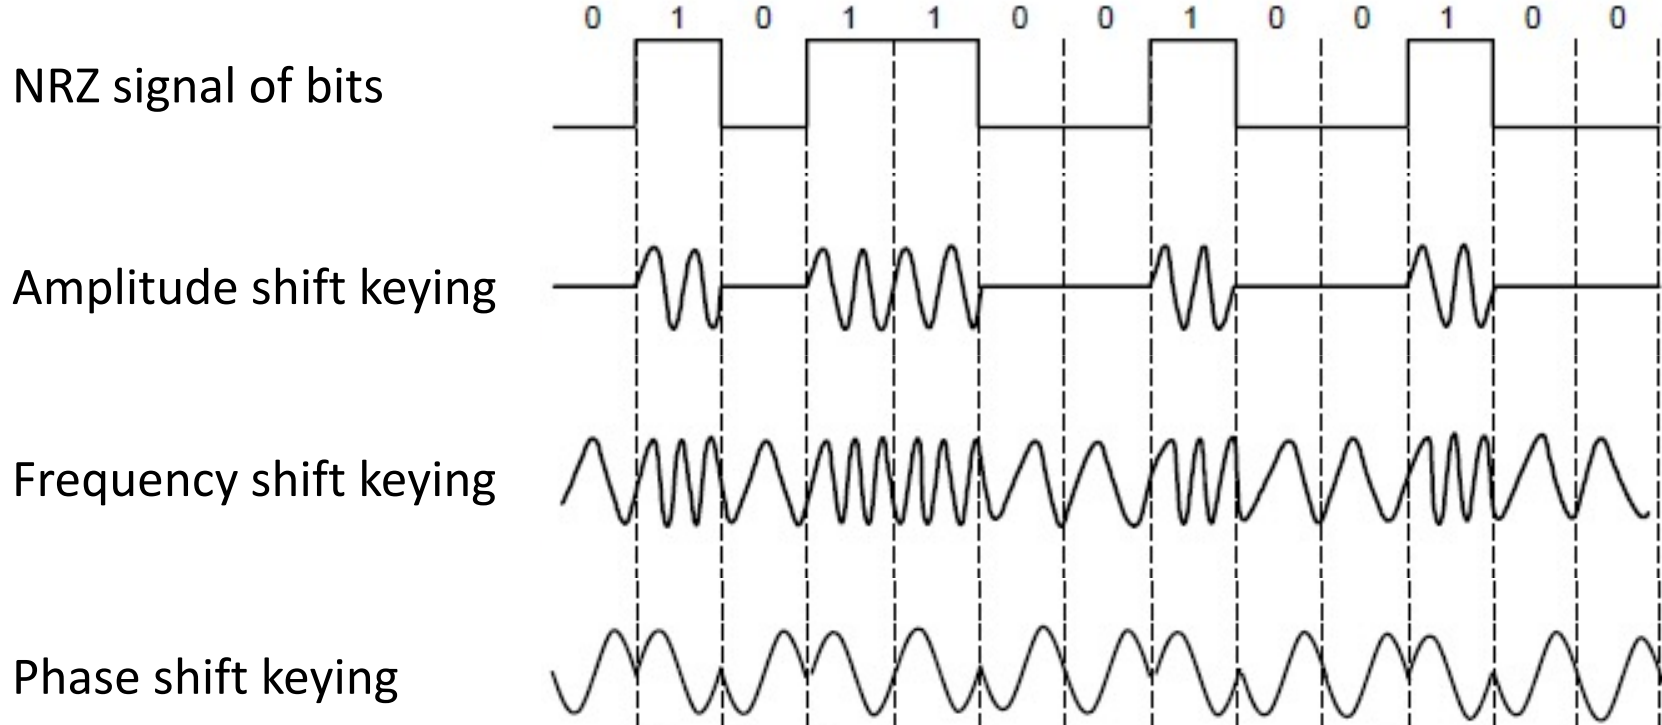


# Sending bits over a link

- We've talked about signals representing bits. How, exactly?
  - This is the topic of modulation




# Modulation



# Key Channel Properties

- The bandwidth (B), signal strength (S), and noise strength (N)
  - B limits the rate of transitions
  - S and N limit how many signal levels we can distinguish



# Claude Shannon (1916-2001)

- Father of information theory
  - “A Mathematical Theory of Communication”, 1948
- Fundamental contributions to digital computers, security, and communications

Electromechanical mouse  
that “solves” mazes! →




Credit: Courtesy MIT Museum

# Shannon Capacity (2)

- Shannon limit is for capacity (C), the maximum information carrying rate of the channel:

$$C = B \log_2(1 + S/BN) \text{ bits/sec}$$



# Wired/Wireless Perspective

- Wires, and Fiber
  - Engineer link to have requisite SNR and B
  - Can fix data rate
- Wireless
  - Given B, but SNR varies greatly, e.g., up to 60 dB!
  - Can't design for worst case, must adapt data rate

# Wired/Wireless Perspective (2)

- Wires, and Fiber

Engineer SNR for data rate

- Engineer link to have requisite SNR and B

- Can fix data rate

- Wireless


Adapt data rate to SNR

- Given B, but SNR varies greatly, e.g., up to 60 dB!

- Can't design for worst case, must adapt data rate

# Putting it all together – DSL

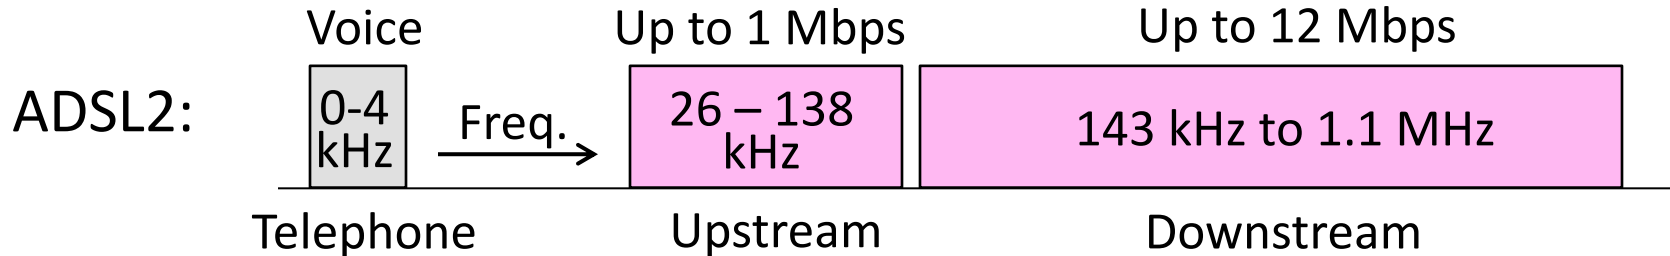
- DSL (Digital Subscriber Line) is widely used for broadband; many variants offer 10s of Mbps
  - Reuses twisted pair telephone line to the home; it has up to ~2 MHz of bandwidth but uses only the lowest ~4 kHz






# DSL (2)

- DSL uses passband modulation (called OFDM)
  - Separate bands for upstream and downstream (larger)
  - Modulation varies both amplitude and phase (called QAM)
  - High SNR, up to 15 bits/symbol, low SNR only 1 bit/symbol




# Error Correction and Detections

- Some bits will be received in error due to noise. What can we do?
  - Detect errors with codes »
  - Correct errors with codes »
  - Retransmit lost frames
- Reliability is a concern that cuts across the layers – we'll see it again




# Approach – Add Redundancy

- Error detection codes
  - Add check bits to the message bits to let some errors be detected
- Error correction codes
  - Add more check bits to let some errors be corrected
- Key issue is now to structure the code to detect many errors with few check bits and modest computation




# Error detection codes

- Some bits may be received in error due to noise. How do we detect this?
  - Parity »
  - Checksums »
  - CRCs »
- Detection will let us fix the error, for example, by retransmission (later).




# Simple Error Detection – Parity Bit

- Take  $D$  data bits, add 1 check bit that is the sum of the  $D$  bits
  - Sum is modulo 2 or XOR



# Parity Bit (2)

- How well does parity work?
  - What is the distance of the code?
  - How many errors will it detect/correct?
- What about larger errors?




# Checksums

- Idea: sum up data in N-bit words
  - Widely used in, e.g., TCP/IP/UDP




- Stronger protection than parity



# Internet Checksum

- Sum is defined in 1s complement arithmetic (must add back carries)
  - And it's the negative sum
- *“The checksum field is the 16 bit one's complement of the one's complement sum of all 16 bit words ...”* – RFC 791





# Internet Checksum (2)

Sending:

1. Arrange data in 16-bit words
2. Put zero in checksum position, add
3. Add any carryover back to get 16 bits
4. Negate (complement) to get sum

0001  
f203  
f4f5  
f6f7

# Internet Checksum (3)

Sending:

1. Arrange data in 16-bit words
2. Put zero in checksum position, add
3. Add any carryover back to get 16 bits
4. Negate (complement) to get sum

```
0001
f203
f4f5
f6f7
+ (0000)
-----
2ddf0
      ↓
 ddf0
+      2
-----
 ddf2
      ↓
 220d
```

# Internet Checksum (4)

Receiving:

1. Arrange data in 16-bit words

2. Checksum will be non-zero, add

3. Add any carryover back to get 16 bits

4. Negate the result and check it is 0

```
0001
f203
f4f5
f6f7
+ 220d
-----
```

# Internet Checksum (5)

Receiving:

1. Arrange data in 16-bit words

2. Checksum will be non-zero, add


3. Add any carryover back to get 16 bits

4. Negate the result and check it is 0

$$\begin{array}{r} 0001 \\ f203 \\ f4f5 \\ f6f7 \\ + 220d \\ \hline 2fffd \\ \phantom{2} \downarrow \\ fffd \\ + \phantom{2} \\ \hline ffff \\ \phantom{ffff} \downarrow \\ \mathbf{0000} \end{array}$$


# Internet Checksum (6)

- How well does the checksum work?
  - What is the distance of the code?
  - How many errors will it detect/correct?
- What about larger errors?




# Cyclic Redundancy Check (CRC)

- Even stronger protection
  - Given  $n$  data bits, generate  $k$  check bits such that the  $n+k$  bits are evenly divisible by a generator  $C$
- Example with numbers:
  - $n = 302$ ,  $k = \text{one digit}$ ,  $C = 3$




# CRCs (2)

- The catch:
  - It's based on mathematics of finite fields, in which “numbers” represent polynomials
  - e.g, 10011010 is  $x^7 + x^4 + x^3 + x^1$
- What this means:
  - We work with binary values and operate using modulo 2 arithmetic



# CRCs (3)

- Send Procedure:
  1. Extend the  $n$  data bits with  $k$  zeros
  2. Divide by the generator value  $C$
  3. Keep remainder, ignore quotient
  4. Adjust  $k$  check bits by remainder
- Receive Procedure:
  1. Divide and check for zero remainder





# CRCs (4)

Data bits:  
1101011111

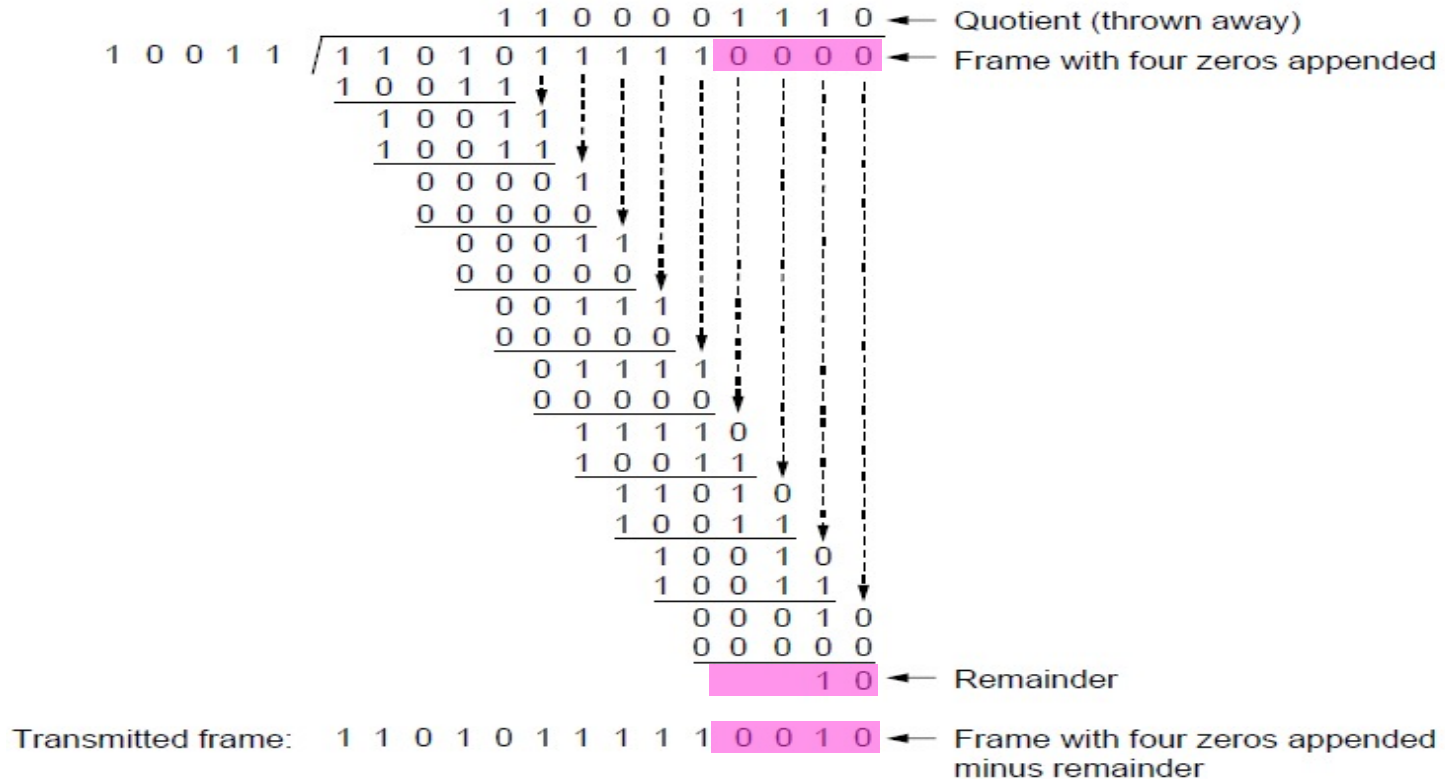
1 0 0 1 1 | 1 1 0 1 0 1 1 1 1 1

Check bits:  
 $C(x) = x^4 + x^1 + 1$

$C = 10011$


$k = 4$

# CRCs (5)




# CRCs (6)

- Protection depend on generator
  - Standard CRC-32 is 10000010  
01100000 10001110 110110111
- Properties:
  - HD=4, detects up to triple bit errors
  - Also odd number of errors
  - And bursts of up to k bits in error
  - Not vulnerable to systematic errors like checksums




# Error Detection in Practice

- CRCs are widely used on links
  - Ethernet, 802.11, ADSL, Cable ...
- Checksum used in Internet
  - TCP, UDP ... but it is weak
- Parity
  - Is little used




# Motivating Example

- A simple code to handle errors:
  - Send two copies! Error if different.
- How good is this code?
  - How many errors can it detect/correct?
  - How many errors will make it fail?




# Motivating Example (2)

- We want to handle more errors with less overhead
  - Will look at better codes; they are applied mathematics
  - But, they can't handle all errors
  - And they focus on accidental errors (will look at secure hashes later)




# Using Error Codes

- Codeword consists of D data plus R check bits (=systematic block code)

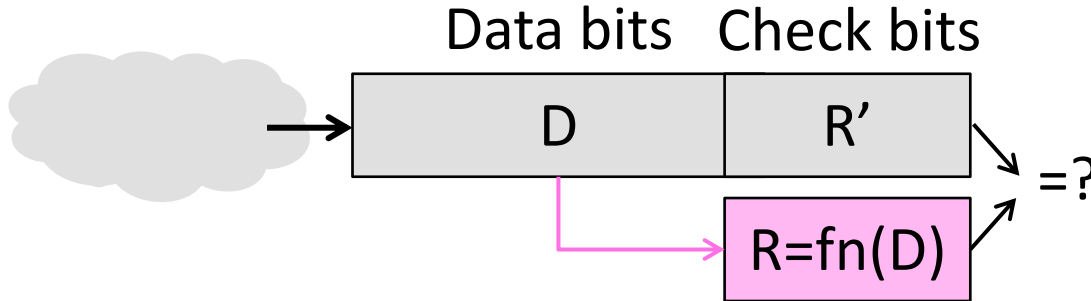


- Sender:
  - Compute R check bits based on the D data bits; send the codeword of D+R bits



# Using Error Codes (2)


- Receiver:
  - Receive  $D+R$  bits with unknown errors
  - Recompute  $R$  check bits based on the  $D$  data bits; error if  $R$  doesn't match  $R'$






# Intuition for Error Codes

- For  $D$  data bits,  $R$  check bits:




- Randomly chosen codeword is unlikely to be correct; overhead is low



# R.W. Hamming (1915-1998)


- Much early work on codes:
  - “Error Detecting and Error Correcting Codes”, BSTJ, 1950
- See also:
  - “You and Your Research”, 1986



Source: IEEE GHN, © 2009 IEEE


# Hamming Distance

- Distance is the number of bit flips needed to change  $D_1$  to  $D_2$
- Hamming distance of a code is the minimum distance between any pair of codewords




# Hamming Distance (2)

- Error detection:
  - For a code of distance  $d+1$ , up to  $d$  errors will always be detected



# Hamming Distance (3)

- Error correction:
  - For a code of distance  $2d+1$ , up to  $d$  errors can always be corrected by mapping to the closest codeword





# ARQ

- ARQ often used when errors are common or must be corrected
  - E.g., WiFi, and TCP (later)
- Rules at sender and receiver:
  - Receiver automatically acknowledges correct frames with an ACK
  - Sender automatically resends after a timeout, until an ACK is received





# ARQ (2)

- Normal operation (no loss)



# ARQ (3)


- Loss and retransmission






# So What's Tricky About ARQ?

- Two non-trivial issues:
  - How long to set the timeout? »
  - How to avoid accepting duplicate frames as new frames »
- Want performance in the common case and correctness always





# Timeouts

- Timeout should be:
  - Not too big (link goes idle)
  - Not too small (spurious resend)
- Fairly easy on a LAN
  - Clear worst case, little variation
- Fairly difficult over the Internet
  - Much variation, no obvious bound
  - We'll revisit this with TCP (later)





# Duplicates

- What happens if an ACK is lost?





# Duplicates (2)

- What happens if an ACK is lost?





# Duplicates (3)

- Or the timeout is early?




# Duplicates (4)

- Or the timeout is early?





# Sequence Numbers

- Frames and ACKs must both carry sequence numbers for correctness
- To distinguish the current frame from the next one, a single bit (two numbers) is sufficient
  - Called Stop-and-Wait



# Stop-and-Wait



- In the normal case:







# Stop-and-Wait (2)

- In the normal case:





# Stop-and-Wait (3)

- With ACK loss:





# Stop-and-Wait (4)

- With ACK loss:





# Stop-and-Wait (5)

- With early timeout:




# Stop-and-Wait (6)

- With early timeout:



# Limitation of Stop-and-Wait


- It allows only a single frame to be outstanding from the sender:
  - Good for LAN, not efficient for high BD




- Ex:  $R=1$  Mbps,  $D = 50$  ms
  - How many frames/sec? If  $R=10$  Mbps?

# Sliding Window

- Generalization of stop-and-wait
  - Allows  $W$  frames to be outstanding
  - Can send  $W$  frames per RTT ( $=2D$ )




- Various options for numbering frames/ACKs and handling loss
  - Will look at along with TCP (later)



# Multiple devices?


- Multiplexing is the network word for the sharing of a resource
- Classic scenario is sharing a link among different users
  - Time Division Multiplexing (TDM) »
  - Frequency Division Multiplexing (FDM) »






# Time Division Multiplexing (TDM)

- Users take turns on a fixed schedule





# Frequency Division Multiplexing (FDM)

- Put different users on different frequency bands





# TDM versus FDM

- In TDM a user sends at a high rate a fraction of the time; in FDM, a user sends at a low rate all the time




# TDM versus FDM (2)

- In TDM a user sends at a high rate a fraction of the time; in FDM, a user sends at a low rate all the time






# TDM/FDM Usage

- Statically divide a resource
  - Suited for continuous traffic, fixed number of users
- Widely used in telecommunications
  - TV and radio stations (FDM)
  - GSM (2G cellular) allocates calls using TDM within FDM





# Multiplexing Network Traffic

- Network traffic is bursty
  - ON/OFF sources
  - Load varies greatly over time




# Multiplexing Network Traffic (2)

- Network traffic is bursty
  - Inefficient to always allocate user their ON needs with TDM/FDM



# Multiplexing Network Traffic (3)


- Multiple access schemes multiplex users according to their demands – for gains of statistical multiplexing






# Multiple Access

- We will look at two kinds of multiple access protocols
  1. Randomized. Nodes randomize their resource access attempts
    - Good for low load situations
  2. Contention-free. Nodes order their resource access attempts
    - Good for high load or guaranteed quality of service situations




# Random MAC

- We will explore random multiple access control (MAC) protocols
  - This is the basis for classic Ethernet
  - Remember: data traffic is bursty




# ALOHA Network

- Seminal computer network connecting the Hawaiian islands in the late 1960s
  - When should nodes send?
  - A new protocol was devised by Norm Abramson ...




# ALOHA Protocol

- Simple idea:
  - Node just sends when it has traffic.
  - If there was a collision (no ACK received) then wait a random time and resend
- That's it!




# ALOHA Protocol (2)

- Some frames will be lost, but many may get through...
- Good idea?





# ALOHA Protocol (3)

- Simple, decentralized protocol that works well under low load!
- Not efficient under high load
  - Analysis shows at most 18% efficiency
  - Improvement: divide time into slots and efficiency goes up to 36%
- We'll look at other improvements



# Classic Ethernet


- ALOHA inspired Bob Metcalfe to invent Ethernet for LANs in 1973
  - Nodes share 10 Mbps coaxial cable
  - Hugely popular in 1980s, 1990s



: © 2009 IEEE

# CSMA (Carrier Sense Multiple Access)



- Improve ALOHA by listening for activity before we send (Doh!)
  - Can do easily with wires, not wireless
- So does this eliminate collisions?
  - Why or why not?







# CSMA (2)

- Still possible to listen and hear nothing when another node is sending because of delay




# CSMA (3)

- CSMA is a good defense against collisions only when BD is small




# CSMA/CD (with Collision Detection)

- Can reduce the cost of collisions by detecting them and aborting (Jam) the rest of the frame time
  - Again, we can do this with wires




# CSMA/CD Complications

- Want everyone who collides to know that it happened
  - Time window in which a node may hear of a collision is  $2D$  seconds



# CSMA “Persistence”



- What should a node do if another node is sending?



- Idea: Wait until it is done, and send

# CSMA “Persistence” (2)

- Problem is that multiple waiting nodes will queue up then collide
  - More load, more of a problem




# CSMA “Persistence” (3)

- Intuition for a better solution
  - If there are  $N$  queued senders, we want each to send next with probability  $1/N$



# Binary Exponential Backoff (BEB)


- Cleverly estimates the probability
  - 1st collision, wait 0 or 1 frame times
  - 2nd collision, wait from 0 to 3 times
  - 3rd collision, wait from 0 to 7 times ...
- BEB doubles interval for each successive collision
  - Quickly gets large enough to work
  - Very efficient in practice







# Classic Ethernet, or IEEE 802.3

- Most popular LAN of the 1980s, 1990s
  - 10 Mbps over shared coaxial cable, with baseband signals
  - Multiple access with “1-persistent CSMA/CD with BEB”




# Modern Ethernet

- Based on switches, not multiple access, but still called Ethernet
  - We'll get to it in a later segment





# Wireless Complications

- Wireless is more complicated than the wired case (Surprise!)
  1. Nodes may have different areas of coverage – doesn't fit Carrier Sense »
  2. Nodes can't hear while sending – can't Collision Detect »




# Different Coverage Areas

- Wireless signal is broadcast and received nearby, where there is sufficient SNR




# Hidden Terminals

- Nodes A and C are hidden terminals when sending to B
  - Can't hear each other (to coordinate) yet collide at B
  - We want to avoid the inefficiency of collisions





# Exposed Terminals

- B and C are exposed terminals when sending to A and D
  - Can hear each other yet don't collide at receivers A and D
  - We want to send concurrently to increase performance



# Nodes Can't Hear While Sending

- With wires, detecting collisions (and aborting) lowers their cost
- More wasted time with wireless



# Possible Solution: MACA

- MACA uses a short handshake instead of CSMA (Karn, 1990)
  - 802.11 uses a refinement of MACA (later)
- Protocol rules:
  1. A sender node transmits a RTS (Request-To-Send, with frame length)
  2. The receiver replies with a CTS (Clear-To-Send, with frame length)
  3. Sender transmits the frame while nodes hearing the CTS stay silent
    - Collisions on the RTS/CTS are still possible, but less likely



# MACA – Hidden Terminals


- $A \rightarrow B$  with hidden terminal C
  1. A sends RTS, to B

A

B


C

D





# MACA – Hidden Terminals (2)

- A → B with hidden terminal C
  2. B sends CTS, to A, and C too





# MACA – Hidden Terminals (3)

- A → B with hidden terminal C
  2. B sends CTS, to A, and C too



# MACA – Hidden Terminals (4)

- A → B with hidden terminal C
  3. A sends frame while C defers



# MACA – Exposed Terminals


- $B \rightarrow A$ ,  $C \rightarrow D$  as exposed terminals
  - B and C send RTS to A and D

A

B


C

D





# MACA – Exposed Terminals (2)

- $B \rightarrow A$ ,  $C \rightarrow D$  as exposed terminals
  - A and D send CTS to B and C





# MACA – Exposed Terminals (3)

- $B \rightarrow A$ ,  $C \rightarrow D$  as exposed terminals
  - A and D send CTS to B and C



# MACA – Exposed Terminals (4)


- $B \rightarrow A$ ,  $C \rightarrow D$  as exposed terminals
  - A and D send CTS to B and C





# 802.11, or WiFi

- Very popular wireless LAN started in the 1990s
- Clients get connectivity from a (wired) AP (Access Point)
- It's a multi-access problem 😊
- Various flavors have been developed over time
  - Faster, more features



# 802.11 Physical Layer

- Uses 20/40 MHz channels on ISM bands
  - 802.11b/g/n on 2.4 GHz
  - 802.11 a/n on 5 GHz
- OFDM modulation (except legacy 802.11b)
  - Different amplitudes/phases for varying SNRs
  - Rates from 6 to 54 Mbps plus error correction
  - 802.11n uses multiple antennas; see “802.11 with Multiple Antennas for Dummies”



# 802.11 CSMA/CA for Multiple Access

- Sender avoids collisions by inserting small random gaps
  - E.g., when both B and C send, C picks a smaller gap, goes first

