


Lectures 15, 16, 17:
Ethernet – 802.3 and 802.11

Nevil Brownlee

314 S2T 2007

Ethernet 9.3

- IEEE 802.3: CSMA/CD on a shared bus



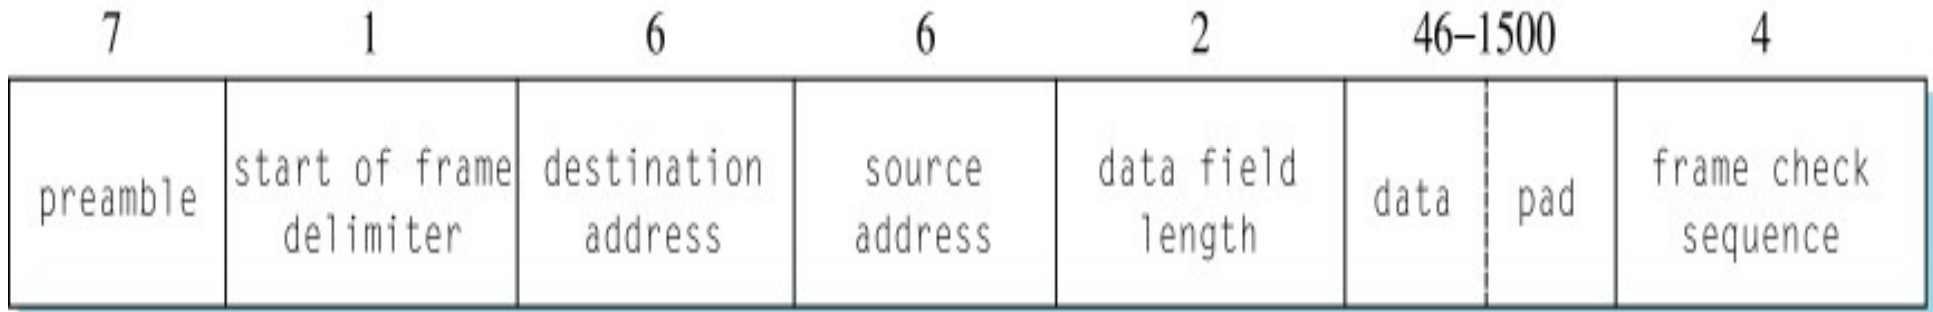
- Transceiver implements the MAC functions
- Originally 10 Mb/s on thick or thin 50Ω cable with repeaters and bridges, later on UTP with hubs and/or switches

Ethernet connection, step by step

- Sending host builds a frame, sends it to Network Interface Card (NIC)
- NIC adds an Ethernet Header, waits for medium idle
- Sends packet, transceiver watches for collision. Tells NIC whether transmission succeeded or failed, NIC retries using *exponential backoff* algorithm
- Receiving host's transceiver sees packet, copies it to its NIC
- That NIC checks packet by computing CRC. If it was for this host (only, or as part of group), sends it to host via interrupt handler

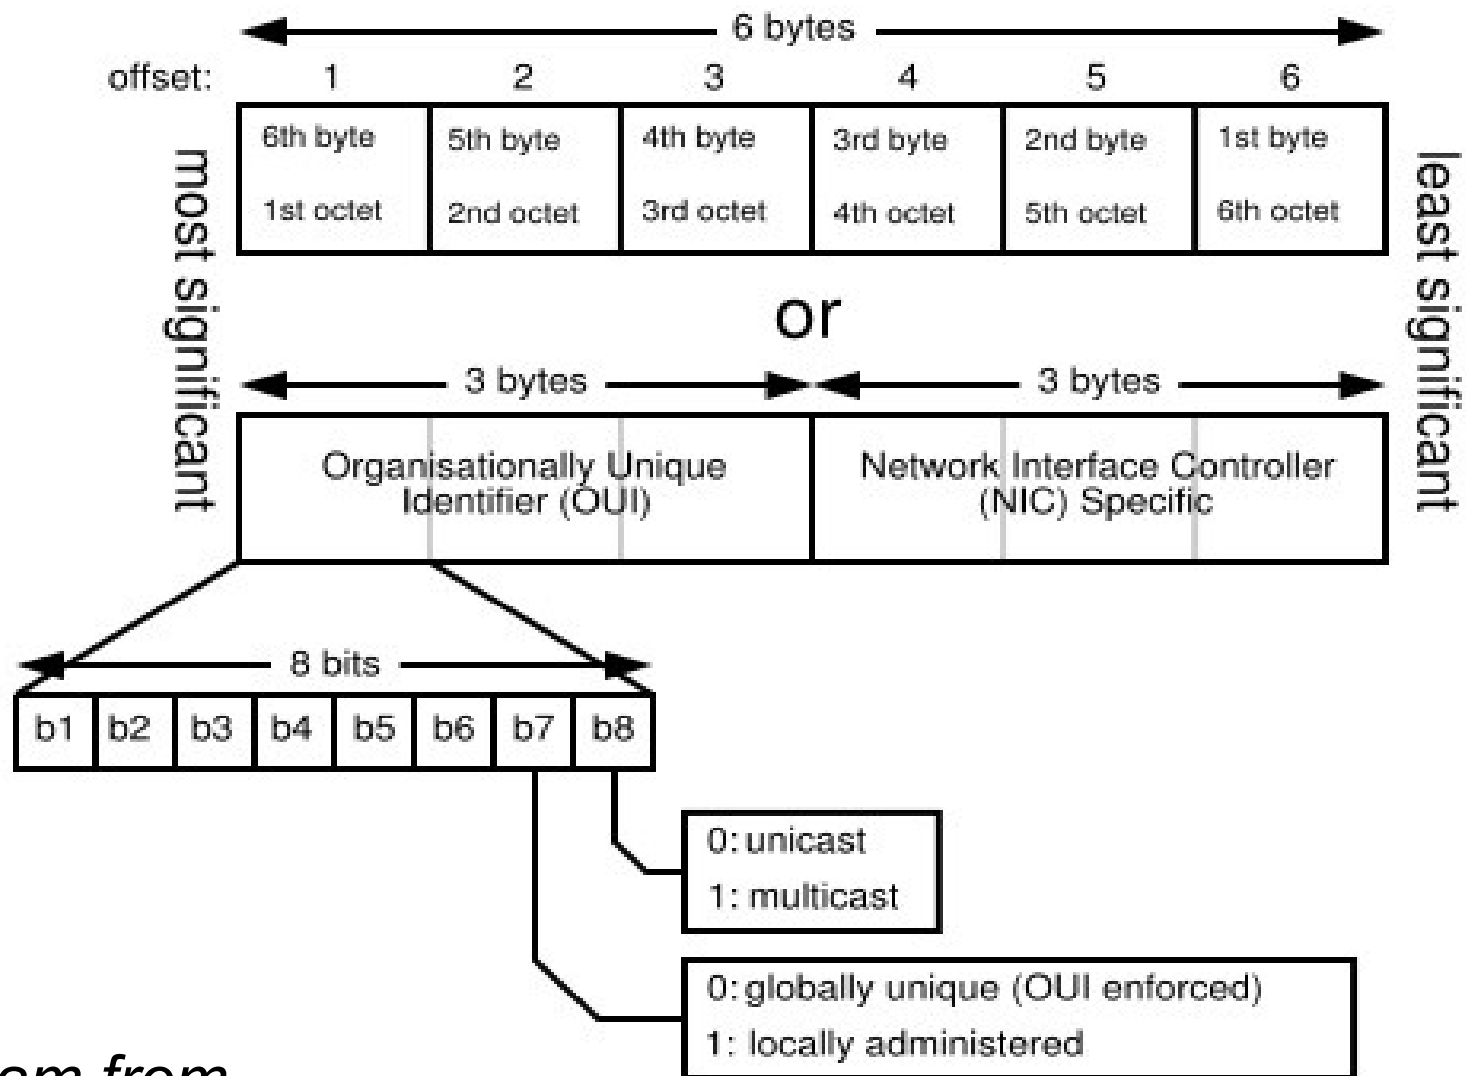
Ethernet Frame, 802.2 encapsulation

number of bytes



- SFD and FCS are not counted as 'packet' bytes – they're not passed in to the host
- Data includes an 802.2 header
- Addresses (6-byte) are globally unique, 48 bits (MAC-48), see next slide
- Ethernet sends bytes in ascending order, bits in a byte low-order-bit-first

Ethernet Address Format (MAC-48)




- *Diagram from Wikipedia web page*

Ethernet Frame, 'native'

- One extra convention:
 - Data Length field can carry an Ethertype instead, provided that the Ethertype value is > 1500 , Ethernet's maximum packet size.
 - For example, Ethertype $0x0800 = 2048$ (IP)
 - Length ≤ 1500 means that an 802.2 header follows

Detecting Collisions

- Max packet size stops a host from monopolising the medium
- Min packet size set for reliable collision detection



10Base5 Ethernet Specifications

- 'Segment' = *collision domain*
- Max segment length **500m**
- Max of **four** repeaters joining segments
- $2500\text{m}/(2 \times 10^8) \text{ m/s} = 12.5 \mu\text{s}$, $25 \mu\text{s}$ round-trip
- Allow $25 \mu\text{s}$ for (worst-case) repeater delay
- $50 \mu\text{s}$ at $10 \text{ Mb/s} = 500 \text{ b}$, plus a few more
- $512 \text{ b} = 64 \text{ B}$

- Min inter-packet gap is **$12.5 \mu\text{s}$** (i.e. 2.5km of cable) for 10, 100 and 1000 Mb/s Ethernet

Physical Implementations


- **10Base5** = Thick Wire
 - thick coax, vampire taps, AUI on (50m) AUI cable
- **10Base2** = Thin Wire
 - thin coax, tee connectors, AUI built into NIC
- **10BaseT** = UTP (unshielded twisted pair) wire
 - max UTP cable length 100 metres
 - UTP into hubs (multiport repeaters) or switches
 - no collisions in switches, allows full-duplex working
 - status pulse to verify link is connected (flashing *link light* on NIC) *[see Wikipedia for details]*

Fast (100 Mb/s) Ethernet 9.4

- 100BaseTX standardised (802.3u) in 1995
- Changes to go from 10 to 100 Mb/s on UTP:
 - couldn't use NRZI encoding directly at 100 Mb/s, too much RF interference (noise)
 - 4B/5B block encoding for each *nibble*, so as to ensure short 'same-bit' runs (Shay Table 9.3)
 - e.g. 1010-0010-0000-0000-0000-0000 becomes 10110-10100-11110-11110-11110-11110
 - that reduced the noise, but not enough to allow use of NRZI
 - MLT-3 signaling ..

Fast (100 Mb/s) Ethernet (2)

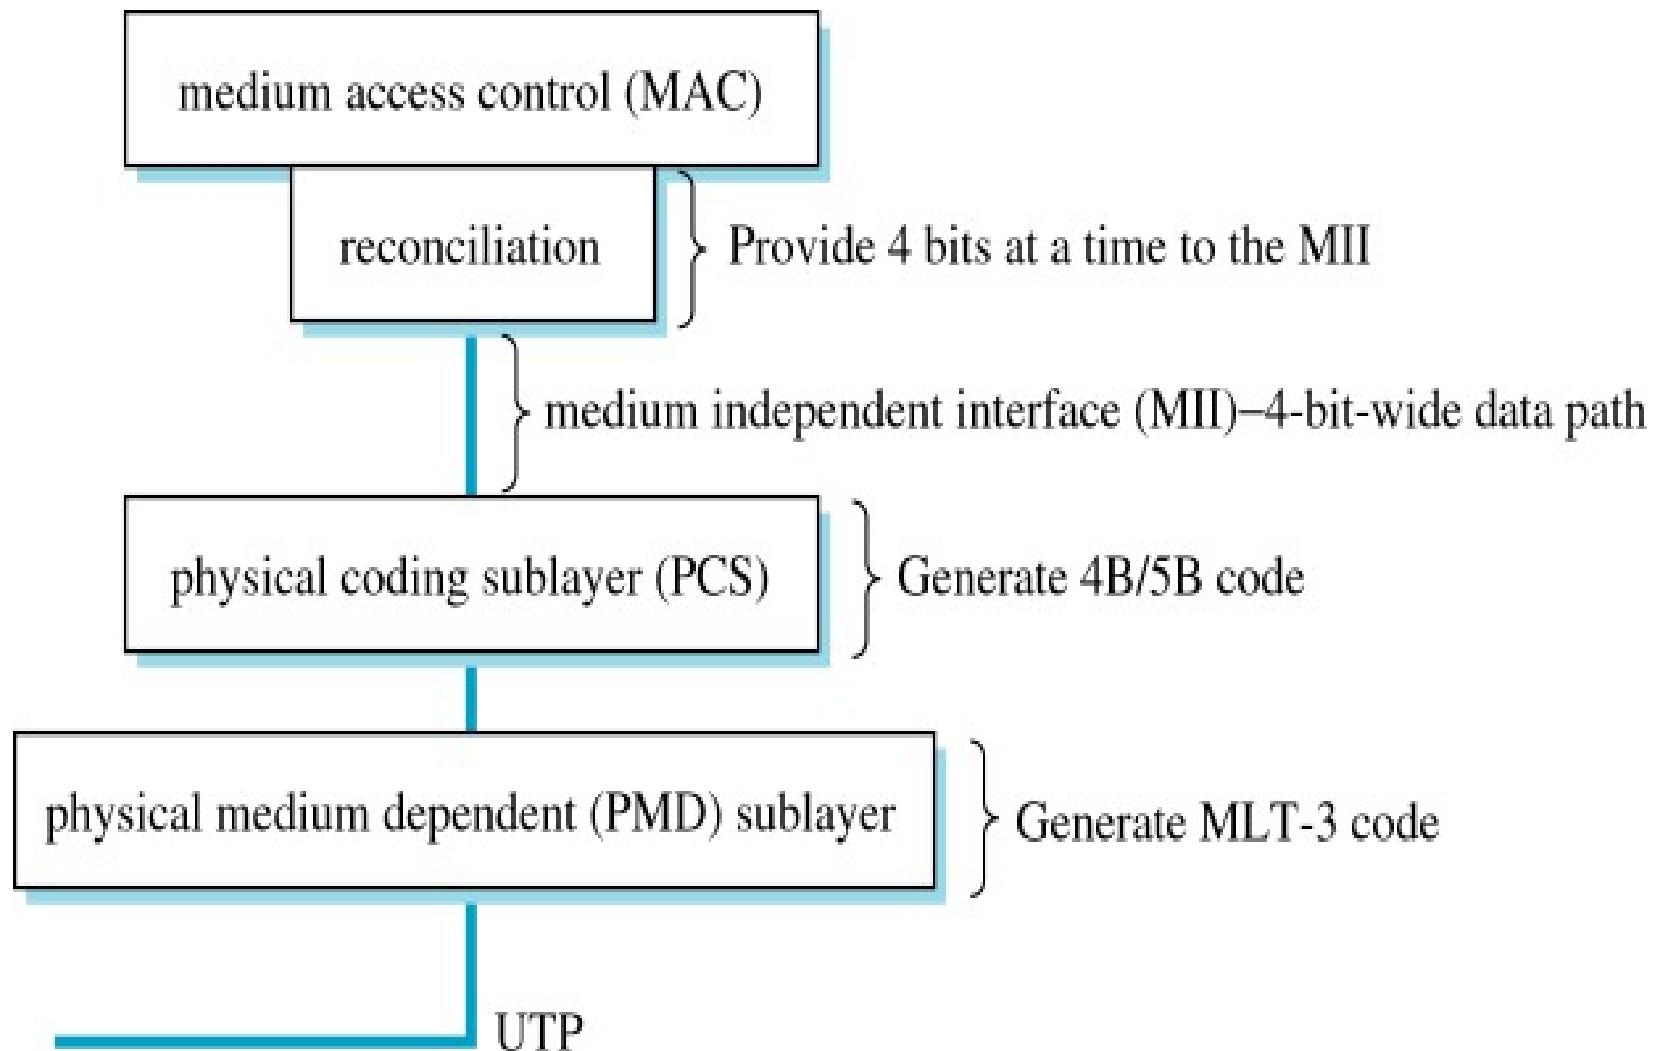
- MLT-3 signaling, Multilevel Line Transmission – Three signal Levels



- MLT-3 cycles through -1, 0, 1, 0, -1, ...
 - for a 1 bit, progress to next state
 - for a 0 bit, maintain same state
- Uses 25% max frequency compared to Manchester, works well over UTP

Fast (100 Mb/s) Ethernet (3)

- 100BaseTX physical layers



100BaseT4

- 100 Mb/s Ethernet on four Category 3 UTP cables
- Not widely used today

100BaseFX – 100 Mb/s on Fibre

- Multi-mode or single-mode fibre
- Segment length 412 metres if collisions can occur, 2 km in full duplex (i.e. using switches)
- Uses 4B/5B block encoding, same as for UTP
- Uses NRZI signaling instead of MLT-3
- Normally use ST fibre connectors
 - ST connectors just push in
 - SC (an older type) is a bayonet-style connector

Collision Domain

- 10Mb/s Ethernet used a minimum frame size of 512 bits, (transmitted in 51.2 μ s) for a maximum segment length of 2500m
- 100Mb/s Ethernet transmits a frame in 1/10 the time, so the max segment length decreases. For 100BaseTX it is only 100m
- 1GB/s Ethernet would require even less!

Gigabit Ethernet 9.5

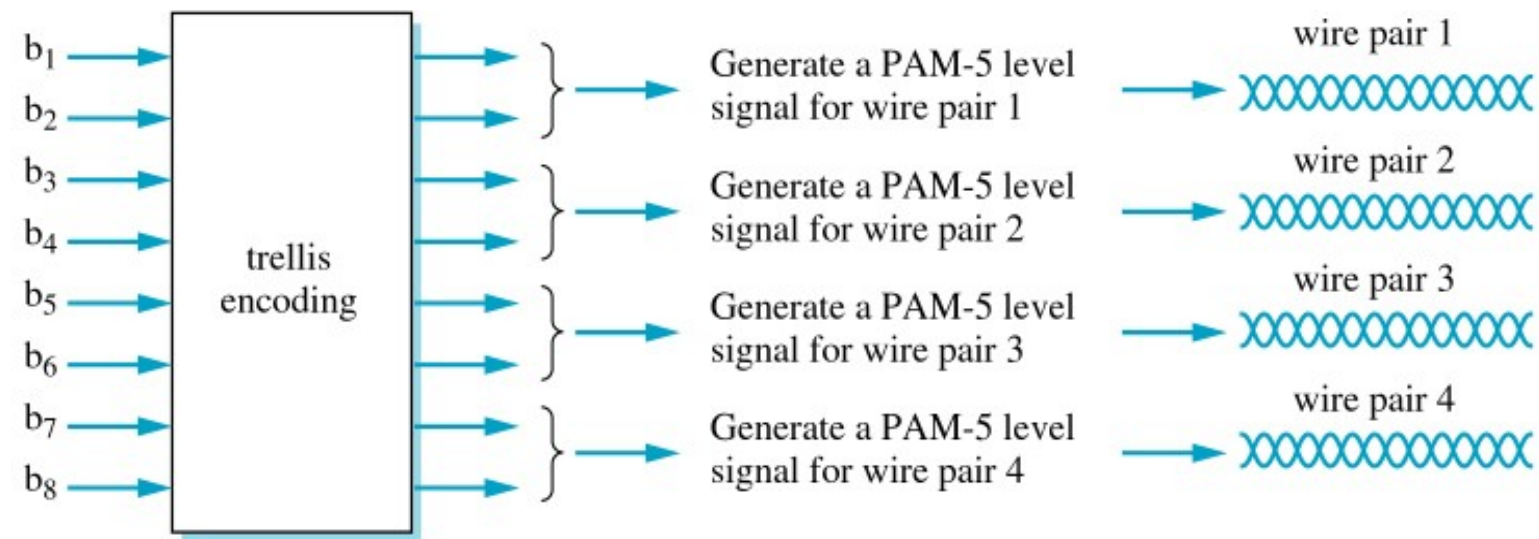
- Collision Domains again ..
 - 1000BaseX (fibre, 802.3z) and 1000BaseT (twisted pair, 802.3ab) allow collisions
 - when collisions are possible, need to use a longer minimum frame so as to keep 100BaseTX's maximum segment length of 100m
 - do that by using a min frame of 4096 bits, i.e. extra padding on short packets
 - can also send a group of packets back-to-back as a 'burst frame,' only the first packet needs to be 4096 bits long
 - collisions are not possible in full-duplex mode; that uses 512b minimum frames (same as earlier standards)

1000BaseX

- Gigabit Ethernet on fibre (or coax cable)
- Similar to 100Mb/s Ethernet, but uses GMII
 - 8-bit-wide data path instead of 4
 - 1 bit of data (on all 8 lines) every 8 ns
- Uses 8B/10B block encoding instead of 4B/5B
 - code symbols are chosen so as to provide *DC balance*, i.e. equal numbers of 0s and 1 over the *long term*
 - has two encoder states and two alternate mappings for each symbol: 'more 0s' and 'more 1s'

1000BaseT

- Gigabit Ethernet over Category 5 UTP
 - Note: 1000BaseTX is a different standard *[not widely used, see Wikipedia]*
- Much harder for UTP than fibre because of its high signal frequencies
- Uses all four twisted pairs in Cat5 cable to carry 250 Mb/s each



From GMII:
8 bits every 8 ns.
1000 Mbps.


One 4D PAM-5 level signal representing 2 bits
on each wire every 8 ns. 250 Mbps per wire pair.

1000BaseT (2)

- 1000BaseT does *not* support half-duplex
- Each GMII octet is divided into four 2-bit groups
- 5-level signalling – PAM5 – is used to send the 2-bit groups. Having 5 levels provides support for some control functions
- Cat5 isn't quite able to carry this reliably, so the link needs error-correction codes to allow for possible errors
 - *trellis encoding* sends extra information, *Viterbi decoding* detects and corrects errors
 - we're not going into the details!

1000BaseT (3)

- All four Cat5 twisted pairs used for data
- Full-duplex carried over each pair at the same time using *hybrids* to combine/separate the signals




10 Gb/s Ethernet

- 802.3ae only works in full-duplex on fibre
- Standard specifies two physical layer types
 - LAN-PHY – for use in LANs
 - e.g. 10GBaseLX4, 300m
 - WAN-PHY – for linking LANs over a wide area
 - e.g. 10GBaseER, 40km
 - an alternative to SONET or ATM

Wireless Networks 9.7

- 802.11 standard; link medium is radio or infrared
- Infrared can bounce off walls and ceiling, radio penetrates through walls
- Normally use one or more *access points* to provide connectivity to movable hosts

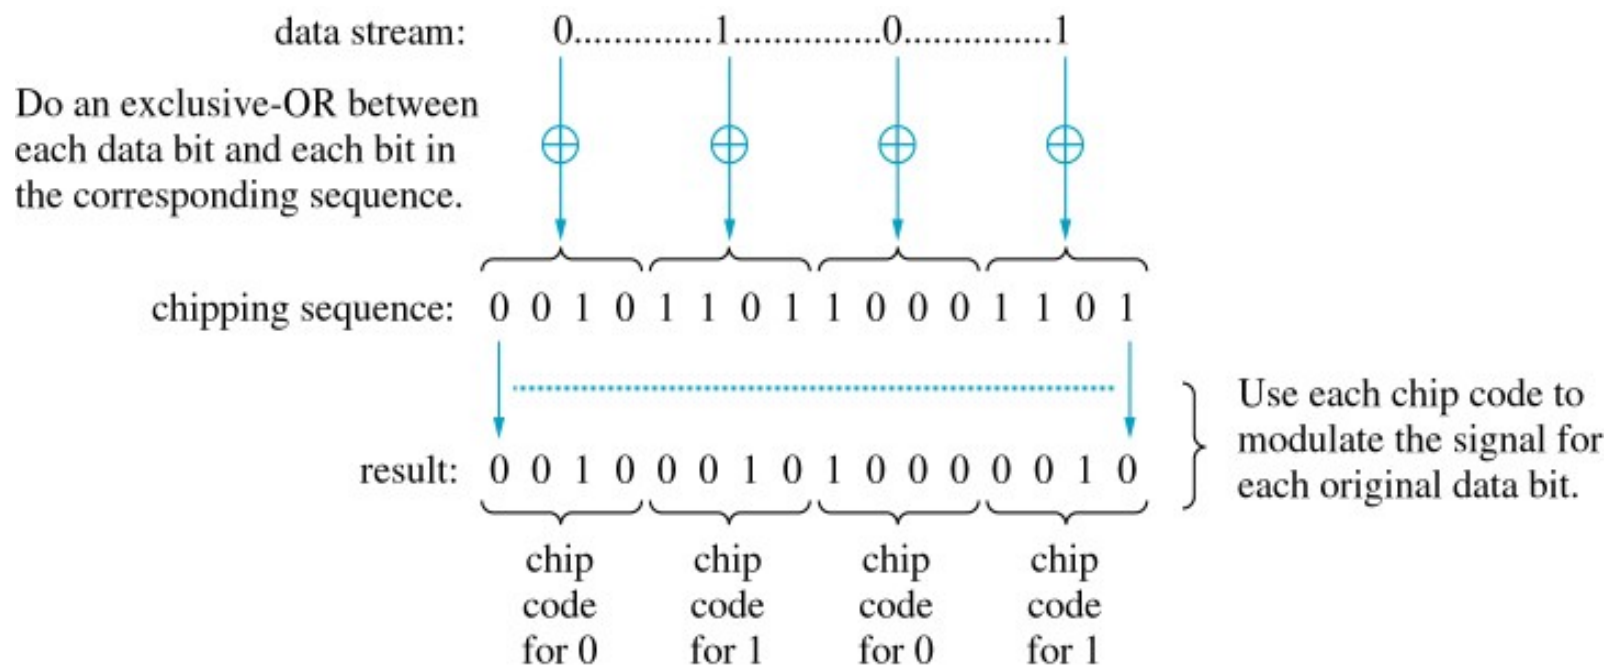


Spread Spectrum Wireless

- Used by 802.11 to minimise interference and (maybe) provide (a little) security
- Two technologies: Frequency Hopping (FHSS) and Direct-Sequence (DSSS)
- FHSS:
 - use a set of frequencies (channels)
 - hop between them in an agreed pseudo-random sequence
 - 802.11 uses 79 channels and 22 hopping sequences

Spread Spectrum Wireless (2)

- DSSS (includes CDMA):
 - for each transmitted bit, send a *chip*, i.e. an n-bit pseudo-random sequence, as illustrated in this diagram



- effect is to generate a high-bandwidth signal, that signal is modulated onto a 2.4 Ghz carrier
- each station uses a different chipping sequence

Contention, *Hidden Station* Problem

- Access Point (AP) can hear all stations, but they can't necessarily hear all of them
- That means they can't detect a collision
- 802.11 has 'Distributed Coordination Function (DCF)' that implements CSMA/CA, i.e. Collision Avoidance
- Next slide illustrates what happens when station A wants to send a message to station B ..

CTS/RTS Protocol

- All devices are contending for the medium. A waits until medium not busy, waits DIFS seconds and sends RTS to B
- B receives RTS and responds with CTS back to A; *however*, it waits for SIFS seconds (a little less than DIFS) before sending. Any other host wanting to send an RTS will wait for DIFS seconds
- If two hosts send RTS at same time the RTS messages will probably collide at B, so B will sense the collision and won't send CTS
- When A receives CTS it knows it has the medium and can send data. When B receives the data it replies with ACK
- Transmission from A to B is now complete, all hosts go back to contending again

802.11 Addressing

- All the stations that communicate with a single AP define a Basic Service Set (BSS)
- BSSes may be connected via a (wired) Distribution System (DS)
- To handle all the stations a host may need to talk to, 802.11 has *four* address fields in its message frame
- Shay lists address field usage in Table 9.9.
 - Address1 for destination host
 - Address2 for source host (needed for sending ACK)

802.11 Frame Format



- Duration: time message will require (for RTS/CTS frames)
- Control: includes ..
 - More Fragments bit. 802.11 may decrease max frame size, fragmenting and reassembling frames as needed. That's done to increase probability of error-free communication
 - To/From DS bit. Set for frames to/from the Distribution System
 - frame Type field. Distinguishes data / control / management frames. RTS, CTS and ACK are control frames

802.11 Management Frames

- Used for:
 - configuring a BSS; *Associate* Request/Response
 - find an AP; *Probe* Req/Resp
 - roaming; *Reassociate* Req/Resp
 - security; *Authenticate* frame, for exchanging security information [keys?]

802.11 Security/Privacy

- WEP – Wired Equivalent Privacy
 - easy to find wireless LANs, e.g. by 'war driving'
 - WEP is a simple authentication/encryption scheme using a 40-bit secret key and a 24-bit initialisation vector. Each message uses a different initialisation vector
 - WEP can be cracked because the initialisation vector sequence may repeat often if traffic is heavy. Newer schemes, e.g. WPA (WI-FI Protected Access) are better
 - 802.11 has several other potential security vulnerabilities

Hints for setting up a WLAN

- You don't *need* to broadcast your SSID (WLAN identifier)
- WEP is better than nothing, use it!
- It's simple to configure an AP to only recognise a small set of 802.11 MAC addresses