Medium Access Control in Wireless Networks

Biplab Sikdar

April 17, 2007

Acknowledgements

Many of the slides are based on a tutorial by Prof. Nitin Vaidya and some other sources.

Overview

- Issues with spectrum sharing
- Current MAC protocols
- Power management
- Rate control
- Interaction with PHY layer

Medium Access Control

- Wireless channel is a shared medium
- MAC coordinates transmission between users sharing the spectrum
- Goals: prevent collisions while maximizing throughput and minimizing delay
- Types:
  - Centralized
  - Decentralized

MAC Protocols: a taxonomy

Three broad classes:

- Channel Partitioning
  - divide channel into smaller "pieces" (time slots, frequency)
  - allocate piece to node for exclusive use
- Random Access
  - allow collisions
  - "recover" from collisions
- "Taking turns"
  - tightly coordinate shared access to avoid collisions

Goal:

Channel Partitioning MAC protocols: TDMA

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

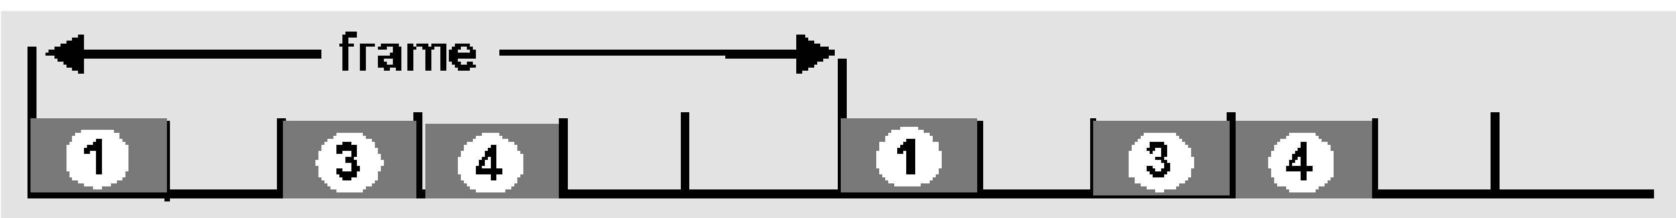

Channel Partitioning MAC protocols: FDMA

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

Random Access Protocols: Unslotted ALOHA

- Simpler, no synchronization
- Packet needs transmission:
  - Send without awaiting for beginning of slot
  - Maximum throughput: 18.4%

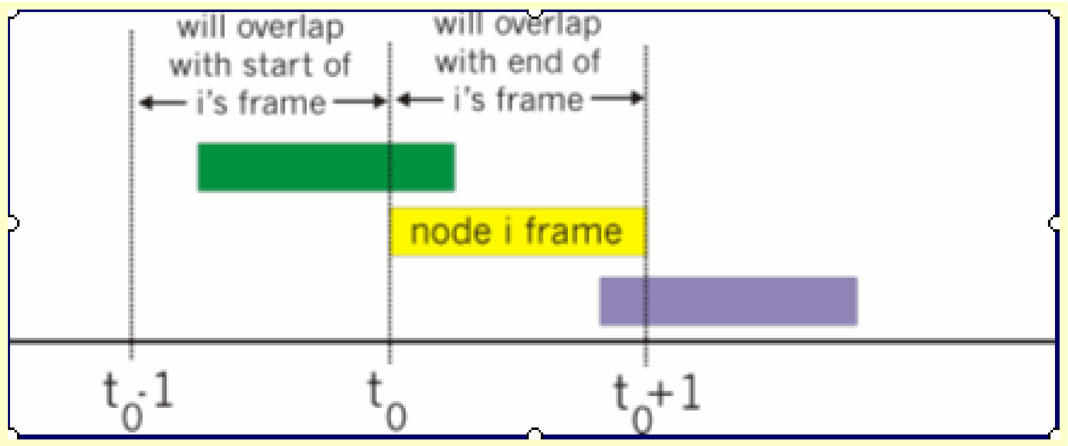

Slotted Aloha

- time is divided into equal size slots (= pkt trans. time)
- node with new arriving pkt: transmit at beginning of next slot
- if collision: retransmit pkt in future slots with probability p, until successful.
- Maximum throughput: 37%

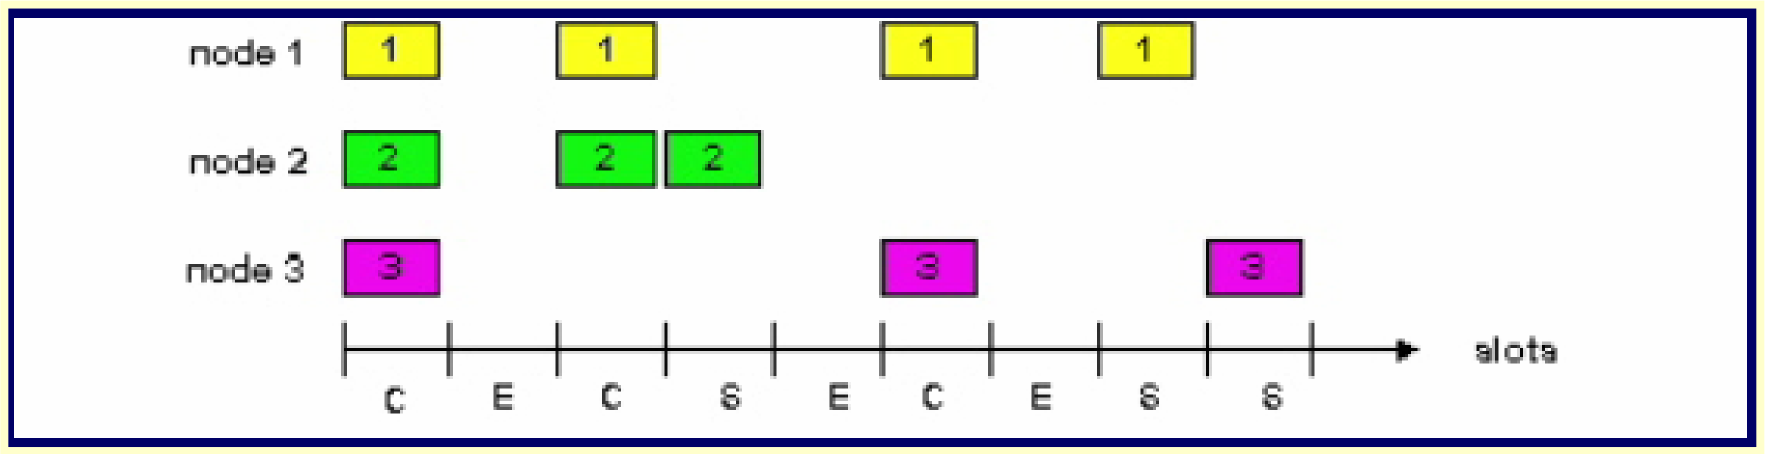

Success (S), Collision (C), Empty (E) slots

Carrier Sense Multiple Access (CSMA)

- In some shorter distance networks, it is possible to listen to the channel before transmitting
- In radio networks, this is called “sensing the carrier”
- The CSMA protocol works just like Aloha except: If the channel is sensed busy, then the user waits to transmit its packet, and a collision is avoided
- This really improves the performance in short distance networks!

Hidden Terminal Problem

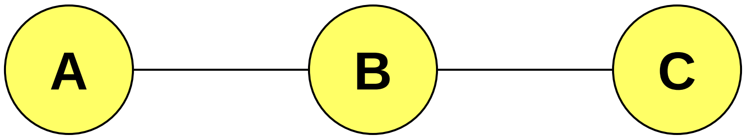

Nodes A and C cannot hear each other

Transmissions by nodes A and C can collide at node B

Nodes A and C are hidden from each other

Busy Tone Solutions [Tobagi75]

- A receiver transmits busy tone when receiving data
- All nodes hearing busy tone keep silent
- Avoids interference from hidden terminals
- Requires a separate channel for busy tone

MACA Solution for Hidden Terminal Problem [Karn90]

- When node A wants to send a packet to node B, node A first sends a *Request-to-Send (RTS)* to A
- On receiving RTS, node A responds by sending *Clear-to-Send (CTS)*, provided node A is able to receive the packet
- When a node (such as C) overhears a CTS, it keeps quiet for the duration of the transfer
  - Transfer duration is included in RTS and CTS both

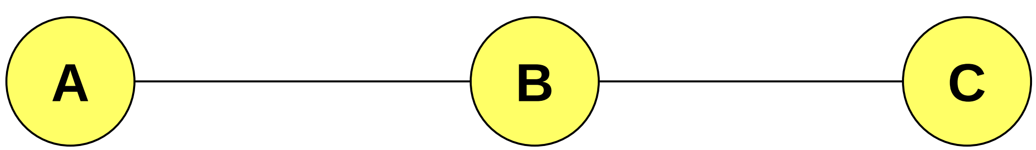

Reliability

- Wireless links are prone to errors. High packet loss rate detrimental to transport-layer performance.
- Mechanisms needed to reduce packet loss rate experienced by upper layers

Simple Solution to Improve Reliability

- When node B receives a data packet from node A, node B sends an Acknowledgement (Ack). This approach adopted in many protocols.
- If node A fails to receive an Ack, it will retransmit the packet.

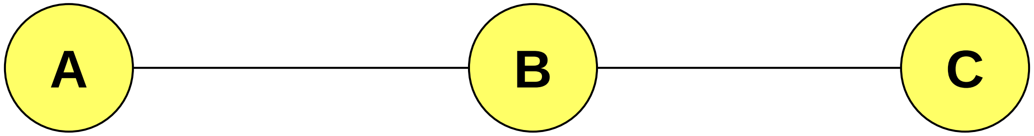

IEEE 802.11 Wireless MAC

- Distributed and centralized MAC components
  - Distributed Coordination Function (DCF)
  - Point Coordination Function (PCF)
- DCF suitable for multi-hop ad hoc networking
- DCF is a Carrier Sense Multiple Access/Collision Avoidance (CSMA/CA) protocol

IEEE 802.11 DCF

- Uses RTS-CTS exchange to avoid hidden terminal problem
  - Any node overhearing a CTS cannot transmit for the duration of the transfer
- Uses ACK to achieve reliability
- Any node receiving the RTS cannot transmit for the duration of the transfer
  - To prevent collision with ACK when it arrives at the sender
  - When B is sending data to C, node A will keep quite

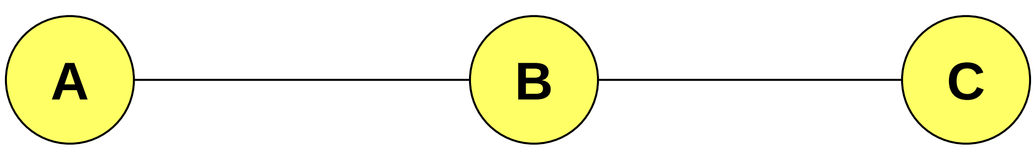

Collision Avoidance

- With half-duplex radios, collision detection is not possible
- **CSMA/CA:** Wireless MAC protocols often use *collision avoidance* techniques, in conjunction with a (physical or virtual) *carrier sense* mechanism
- Carrier sense: When a node wishes to transmit a packet, it first waits until the channel is idle.
- Collision avoidance: Nodes hearing RTS or CTS stay silent for the duration of the corresponding transmission. Once channel becomes idle, the node waits for a randomly chosen duration before attempting to transmit.

IEEE 802.11

RTS = Request-to-Send

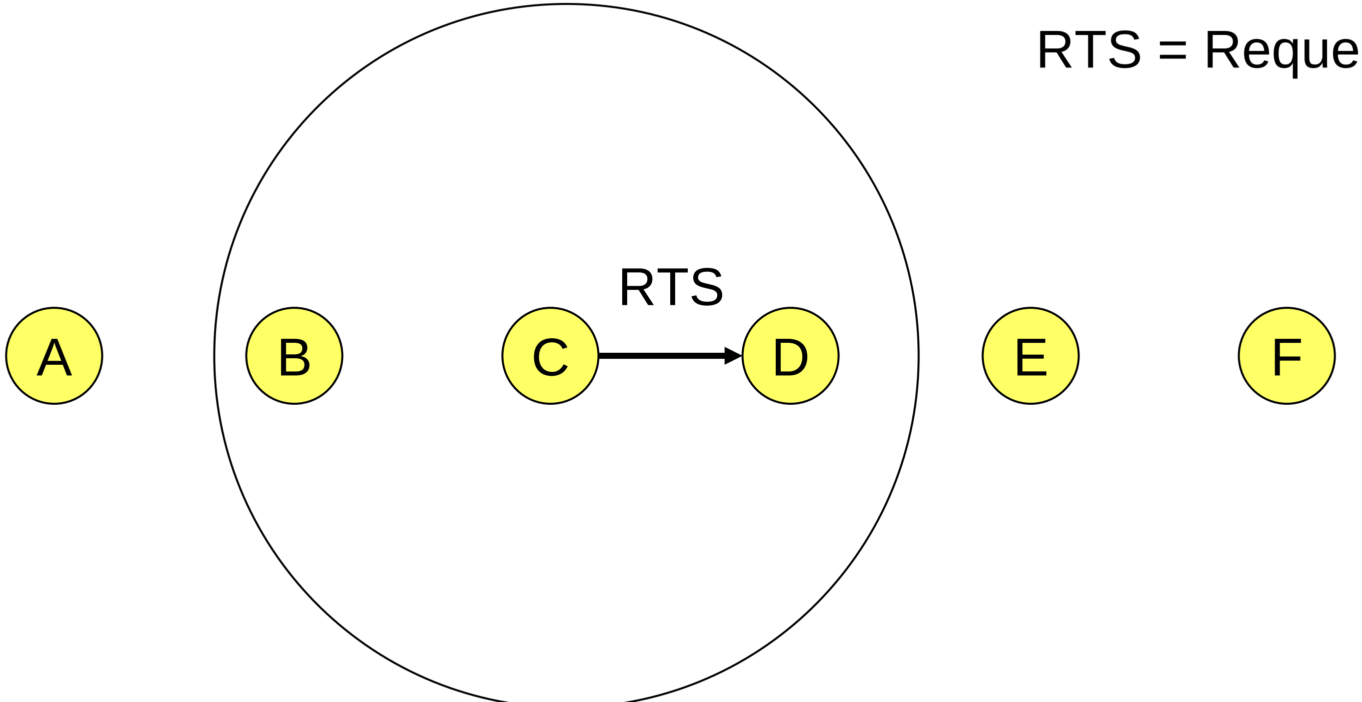

Pretending a circular range

IEEE 802.11

RTS = Request-to-Send

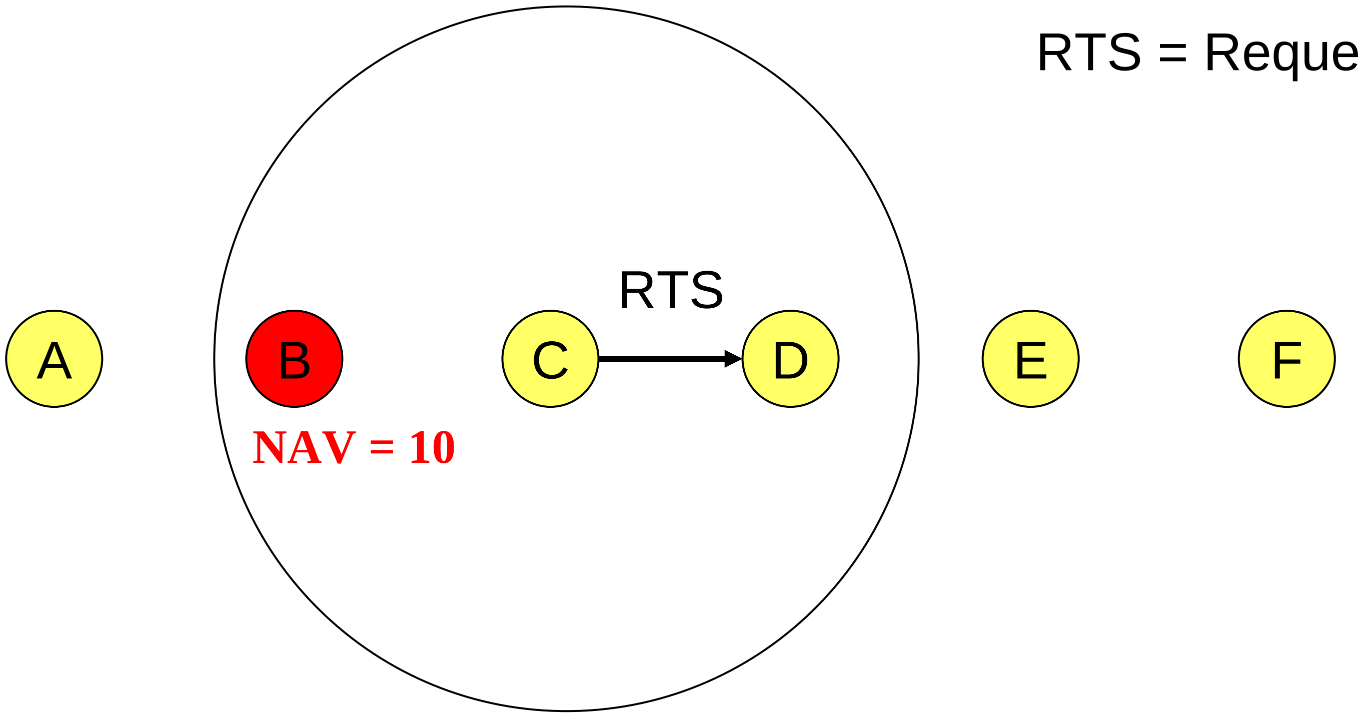

NAV = remaining duration to keep quiet

IEEE 802.11

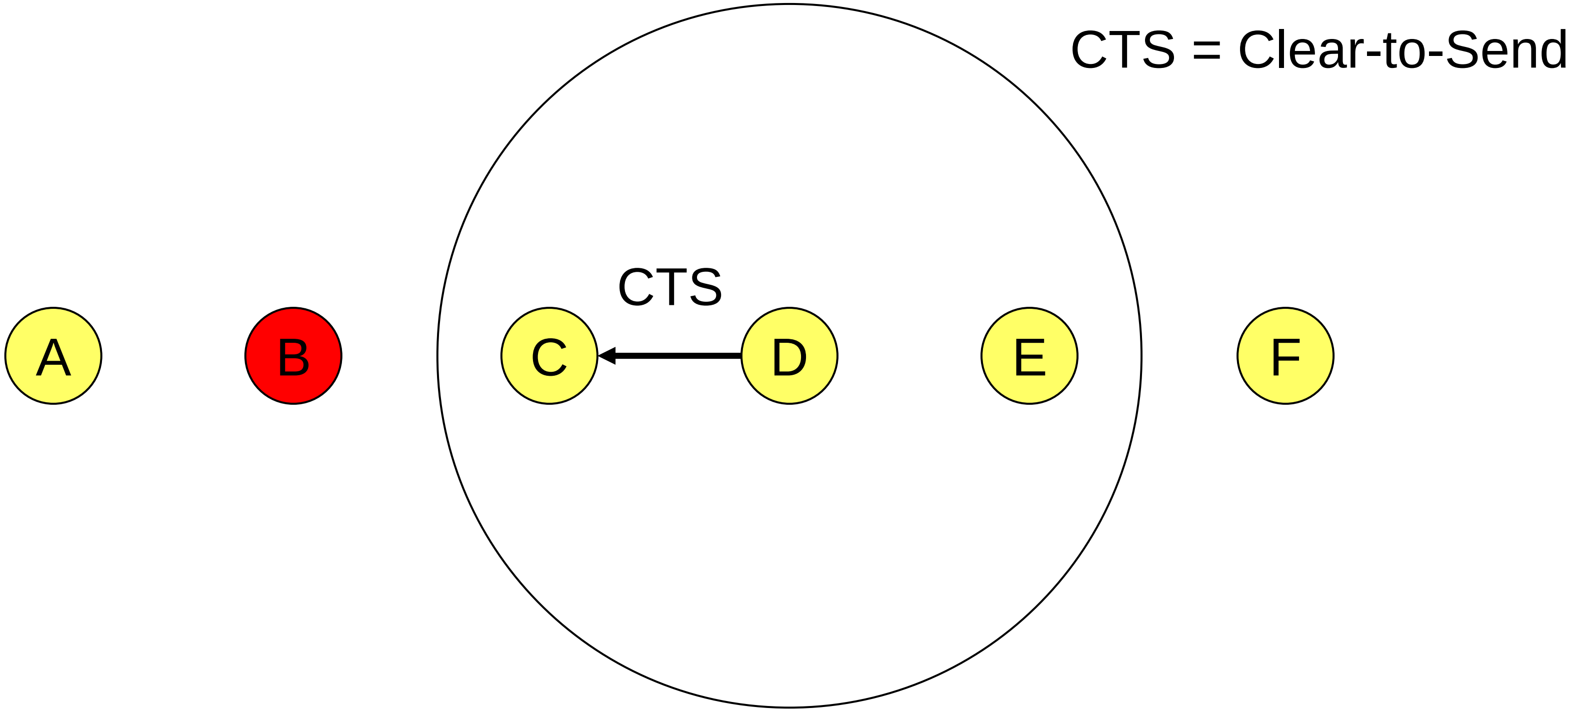

IEEE 802.11

CTS = Clear-to-Send

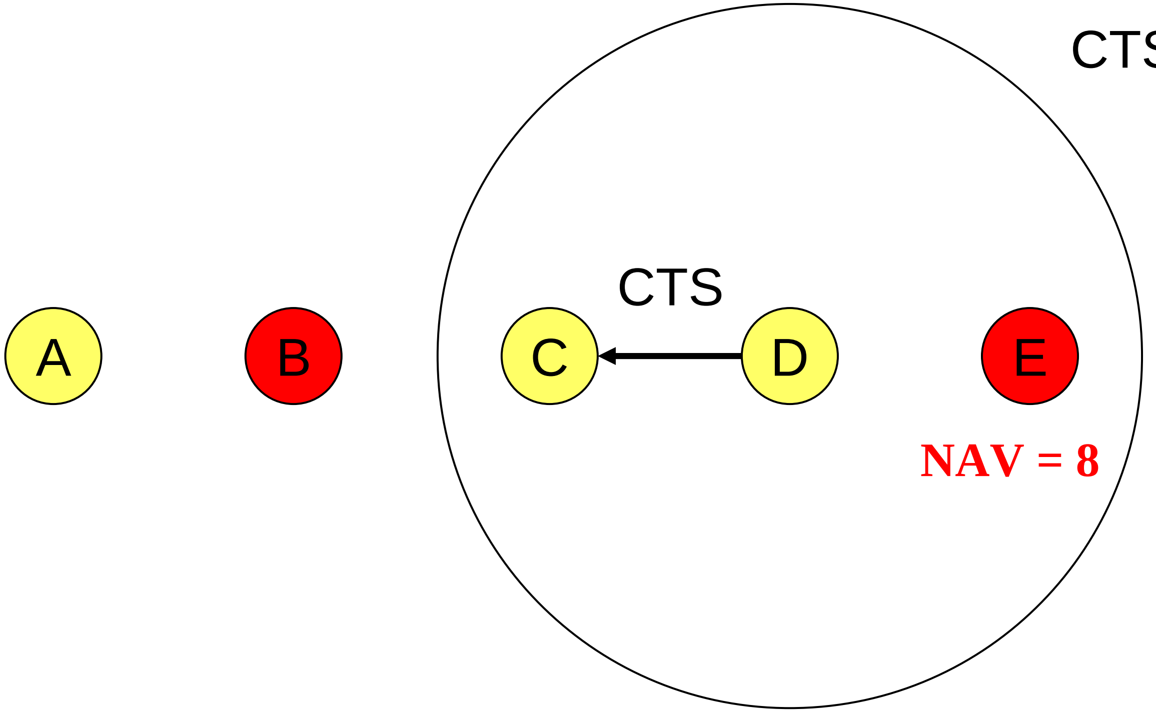

IEEE 802.11

- DATA packet follows CTS. Successful data reception acknowledged using ACK.

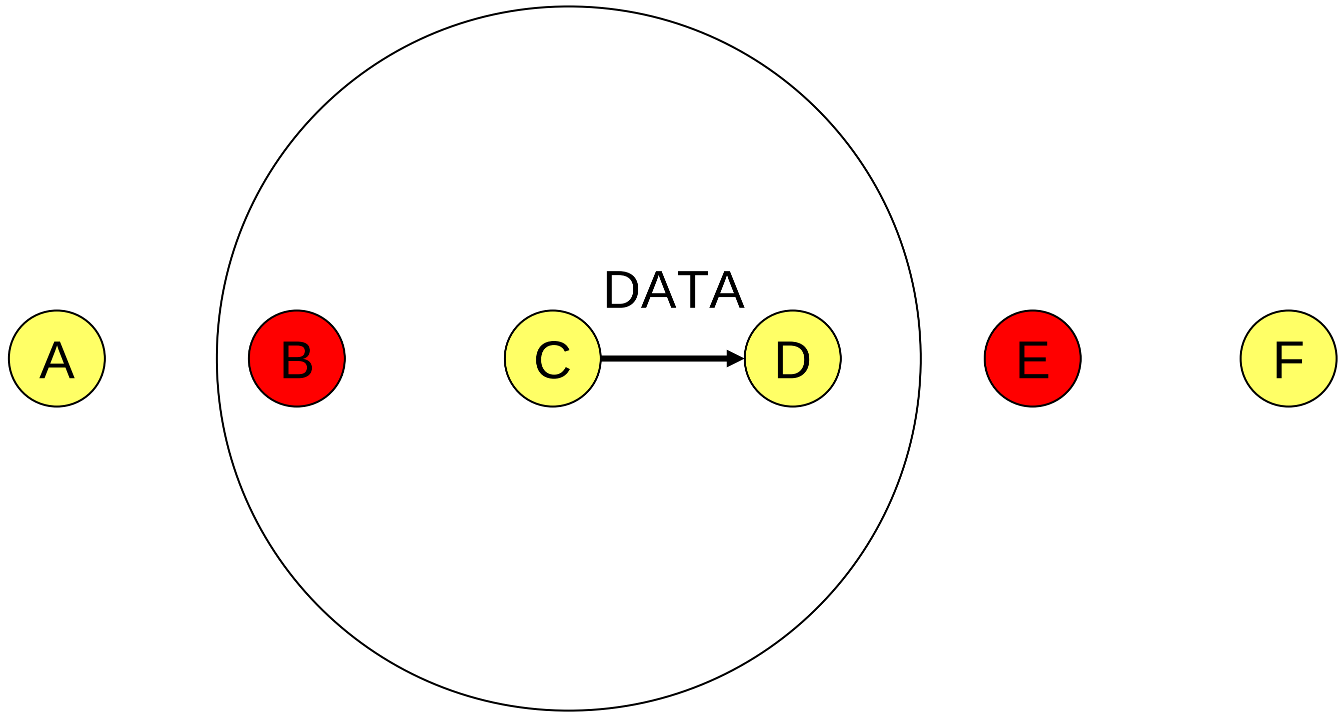

IEEE 802.11

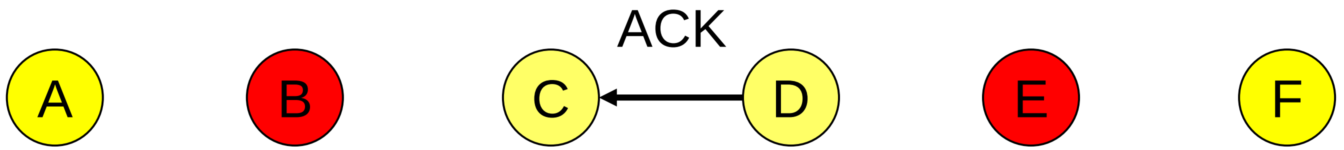

IEEE 802.11

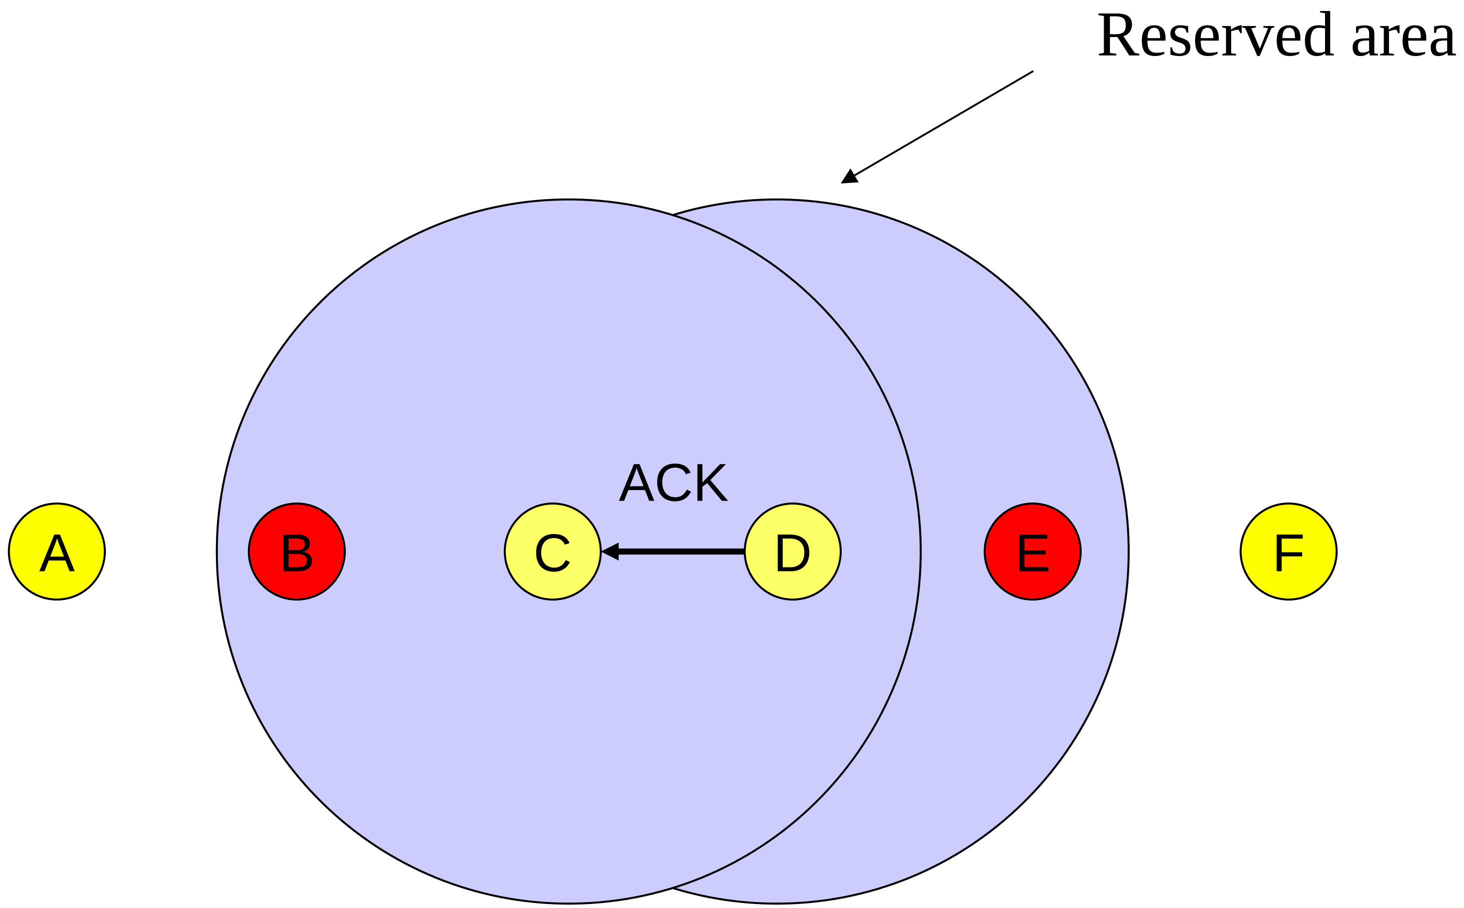

IEEE 802.11

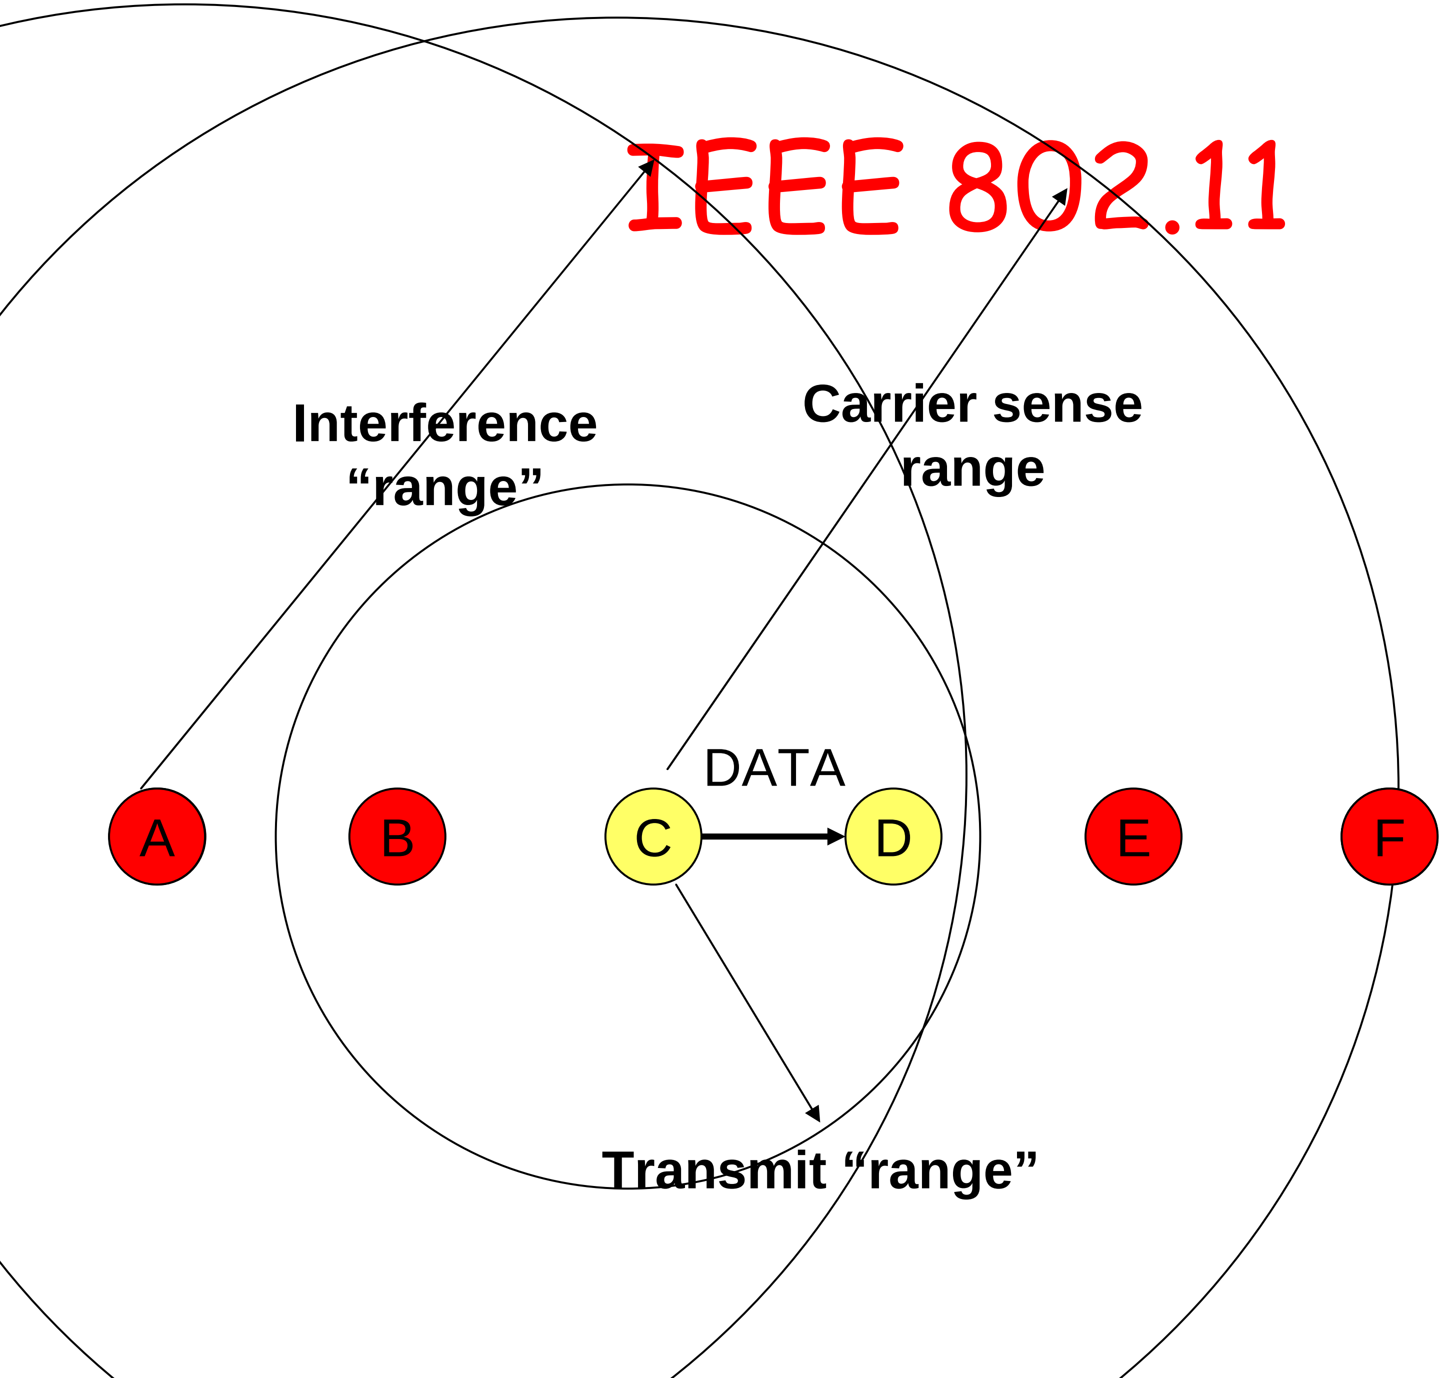
Interference "range"
Carrier sense range
DATA
A
B
C
D
E
F
Transmit "range"

CSMA/CA

- Physical carrier sense, and
- Virtual carrier sense using Network Allocation Vector (NAV)
- NAV is updated based on overheard RTS/CTS/DATA/ACK packets, each of which specified duration of a pending transmission
- Nodes stay silent when carrier sensed (physical/virtual)
- *Backoff intervals* used to reduce collision probability

Backoff Interval

- When transmitting a packet, choose a backoff interval in the range [0,cw]
  - cw is contention window
- Count down the backoff interval when medium is idle
  - Count-down is suspended if medium becomes busy
- When backoff interval reaches 0, transmit RTS

DCF Example

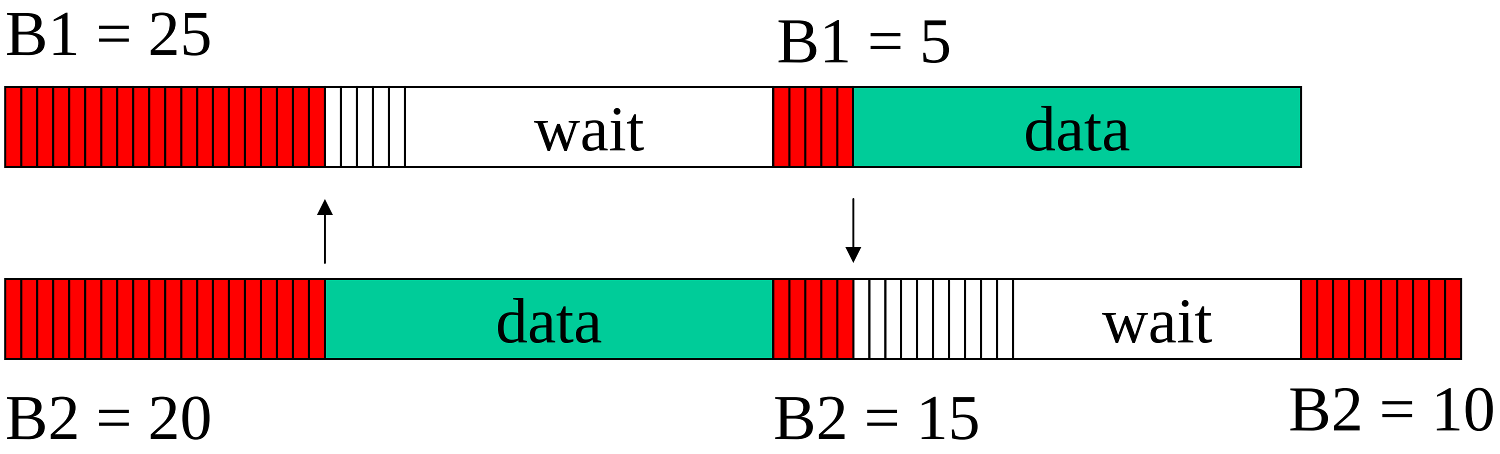

cw = 31

B1 and B2 are backoff intervals at nodes 1 and 2

Backoff Interval

- The time spent counting down backoff intervals is a part of MAC overhead
- Choosing a *large cw* leads to large backoff intervals and can result in larger overhead
- Choosing a *small cw* leads to a larger number of collisions (when two nodes count down to 0 simultaneously)

Backoff Interval

- Since the number of nodes attempting to transmit simultaneously may change with time, some mechanism to manage contention is needed
- IEEE 802.11 DCF: contention window *cw* is chosen dynamically depending on collision occurrence

Binary Exponential Backoff in DCF

- When a node fails to receive CTS in response to its RTS, it increases the contention window
  - *cw* is doubled (up to an upper bound)
- When a node successfully completes a data transfer, it restores *cw* to *CWmin*
- *cw* follows a sawtooth curve

Fairness Issues in MAC

- Many definitions of fairness plausible
- Simplest definition: All nodes should receive *equal* bandwidth

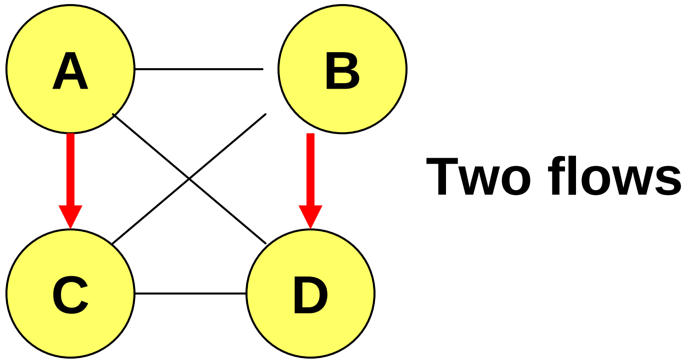

Fairness Issues in MAC

- Assume that initially, A and B both choose a backoff interval in range [0,31] but their RTSs collide
- Nodes A and B then choose from range [0,63]
  - Node A chooses 4 slots and B choose 60 slots
  - After A transmits a packet, it next chooses from range [0,31]
  - It is possible that A may transmit several packets before B transmits its first packet

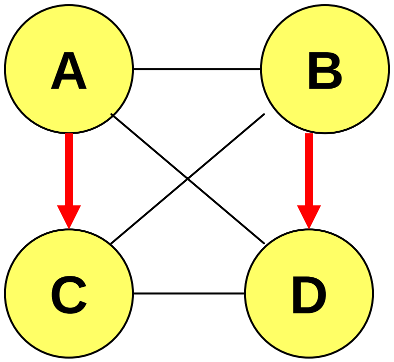

Two flows

Fairness Issues in MAC

- Unfairness occurs when one node has backed off much more than some other node

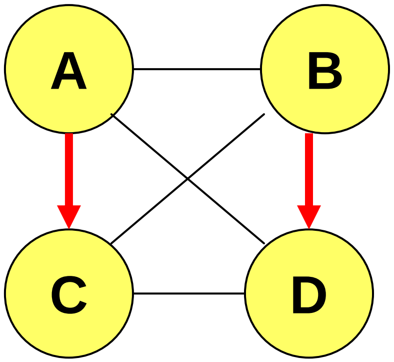

Two flows

MACAW Solution for Fairness [Bharganav94]

- When a node transmits a packet, it appends the *cw* value to the packet, all nodes hearing that *cw* value use it for their future transmission attempts
- Since *cw* is an indication of the level of congestion in the vicinity of a specific receiver node, MACAW proposes maintaining *cw* independently for each receiver
- Using per-receiver *cw* is particularly useful in multi-hop environments, since congestion level at different receivers can be very different

Another MACAW Proposal

- For the scenario below, when node A sends an RTS to B, while node C is receiving from D, node B cannot reply with a CTS, since B knows that D is sending to C
- When the transfer from C to D is complete, node B can send a Request-to-send-RTS to node A
  - Node A may then immediately send RTS to node B

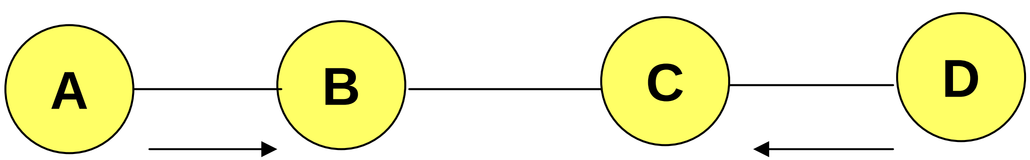

Problems

- This approach, however, does not work in the scenario below
  - Node B may not receive the RTS from A at all, due to interference with transmission from C

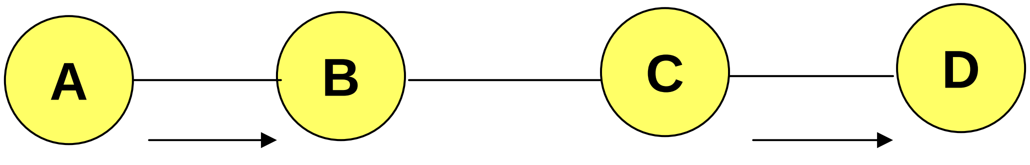

Energy Conservation

- Since many mobile hosts are operated by batteries, MAC protocols which conserve energy are of interest
- Two approaches to reduce energy consumption
  - Power save: Turn off wireless interface when desirable
  - Power control: Reduce transmit power

Power Aware Multi-Access Protocol (PAMAS) [Singh98]

- A node powers off its radio while a neighbor is transmitting to someone else

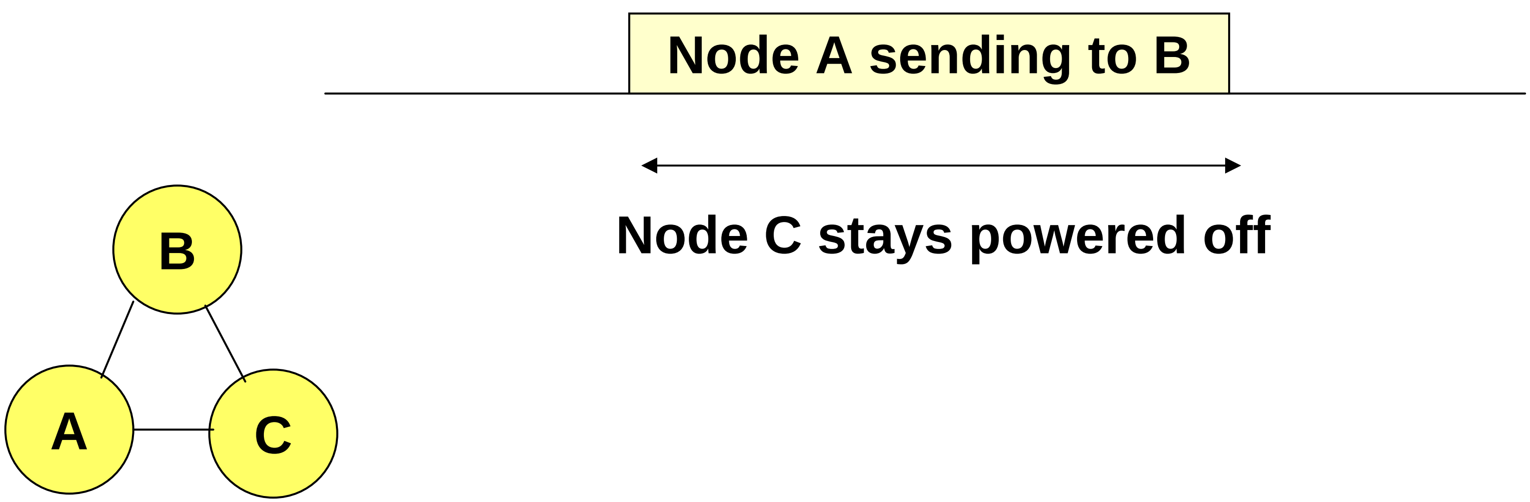

Power Aware Multi-Access Protocol (PAMAS)

- What should node C do when it wakes up and finds that D is transmitting to someone else
  - C does not know how long the transfer will last

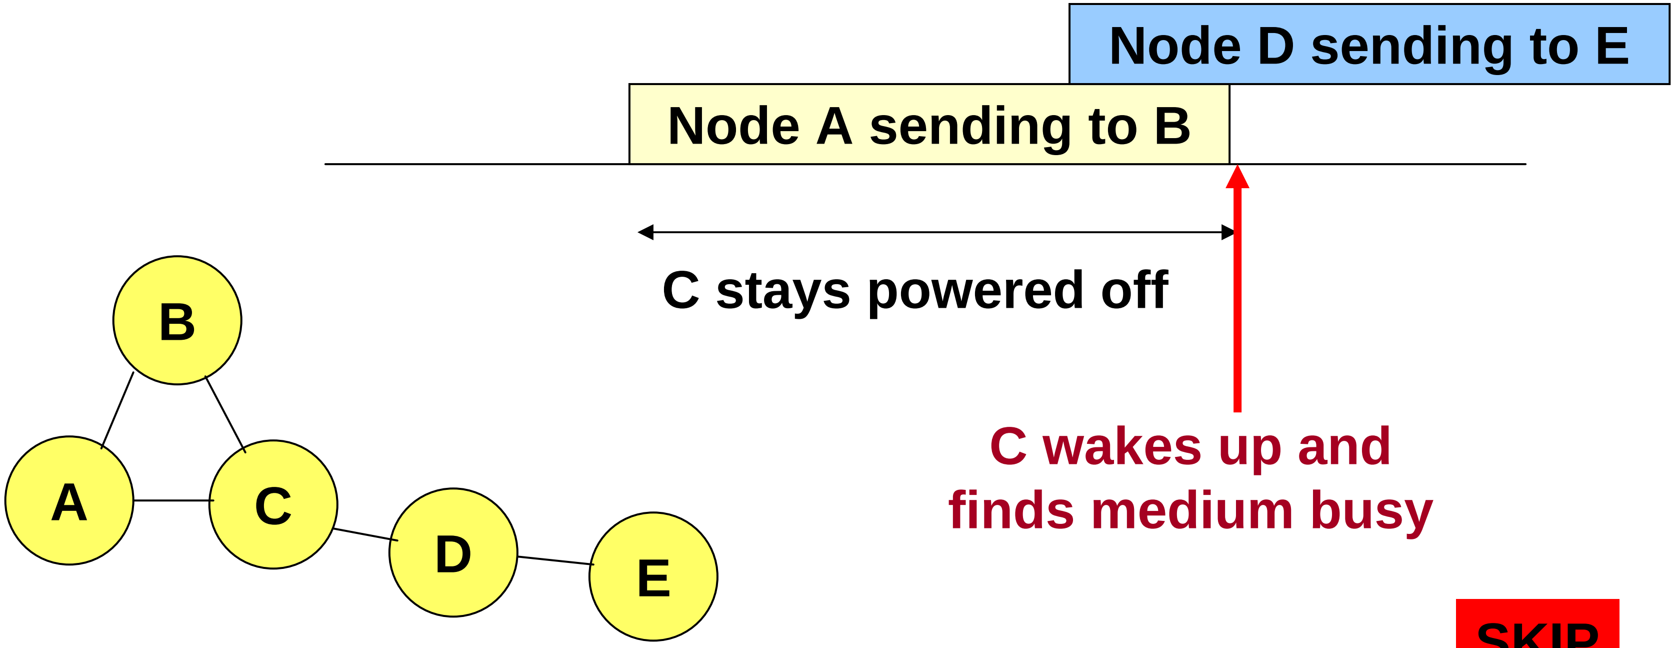

SKIP

PAMAS

- PAMAS uses a control channel separate from the data channel
- Node C on waking up performs a binary probe to determine the length of the longest remaining transfer
  - C sends a probe packet with parameter L
  - All nodes which will finish transfer in interval [L/2,L] respond
  - Depending on whether node C sees silence, collision, or a unique response it takes varying actions
- Node C (using procedure above) determines the duration of time to go back to sleep

Disadvantages of PAMAS

- Use of a separate control channel
- Nodes have to be able to receive on the control channel while they are transmitting on the data channel
  - And also transmit on data and control channels simultaneously
- A node (such as C) should be able to determine when probe responses from multiple senders collide

Another Proposal in PAMAS

- To avoid the probing, a node should switch off the interface for data channel, but not for the control channel (which carries RTS/CTS packets)
- Advantage: Each sleeping node always knows how long to sleep by watching the control channel
- Disadvantage: This may not be useful when hardware is shared for the control and data channels
  - It may not be possible turn off much hardware due to the sharing

Power Save in IEEE 802.11 Ad Hoc Mode

- Time is divided into beacon intervals

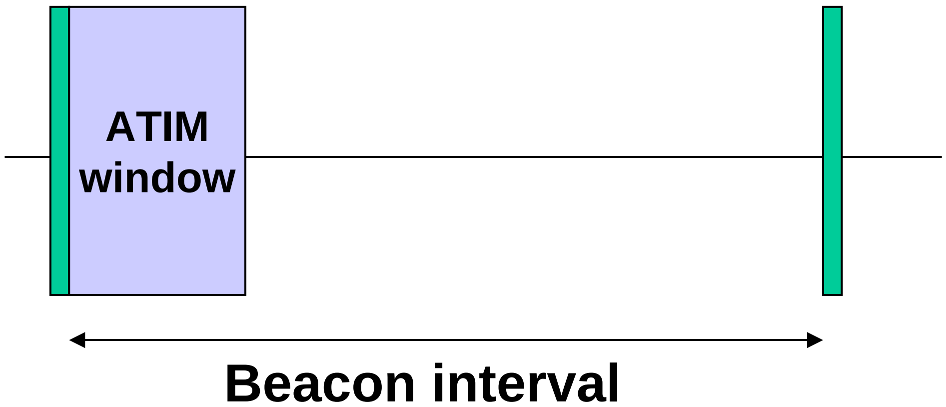

- Each beacon interval begins with an ATIM window

Power Save in IEEE 802.11 Ad Hoc Mode

- If host A has a packet to transmit to B, A must send an ATIM Request to B during an ATIM Window
- On receipt of ATIM Request from A, B will reply by sending an ATIM Ack, and stay up during the rest of the beacon interval
- If a host does not receive an ATIM Request during an ATIM window, and has no pending packets to transmit, it may sleep during rest of the beacon interval

Power Save in IEEE 802.11 Ad Hoc Mode

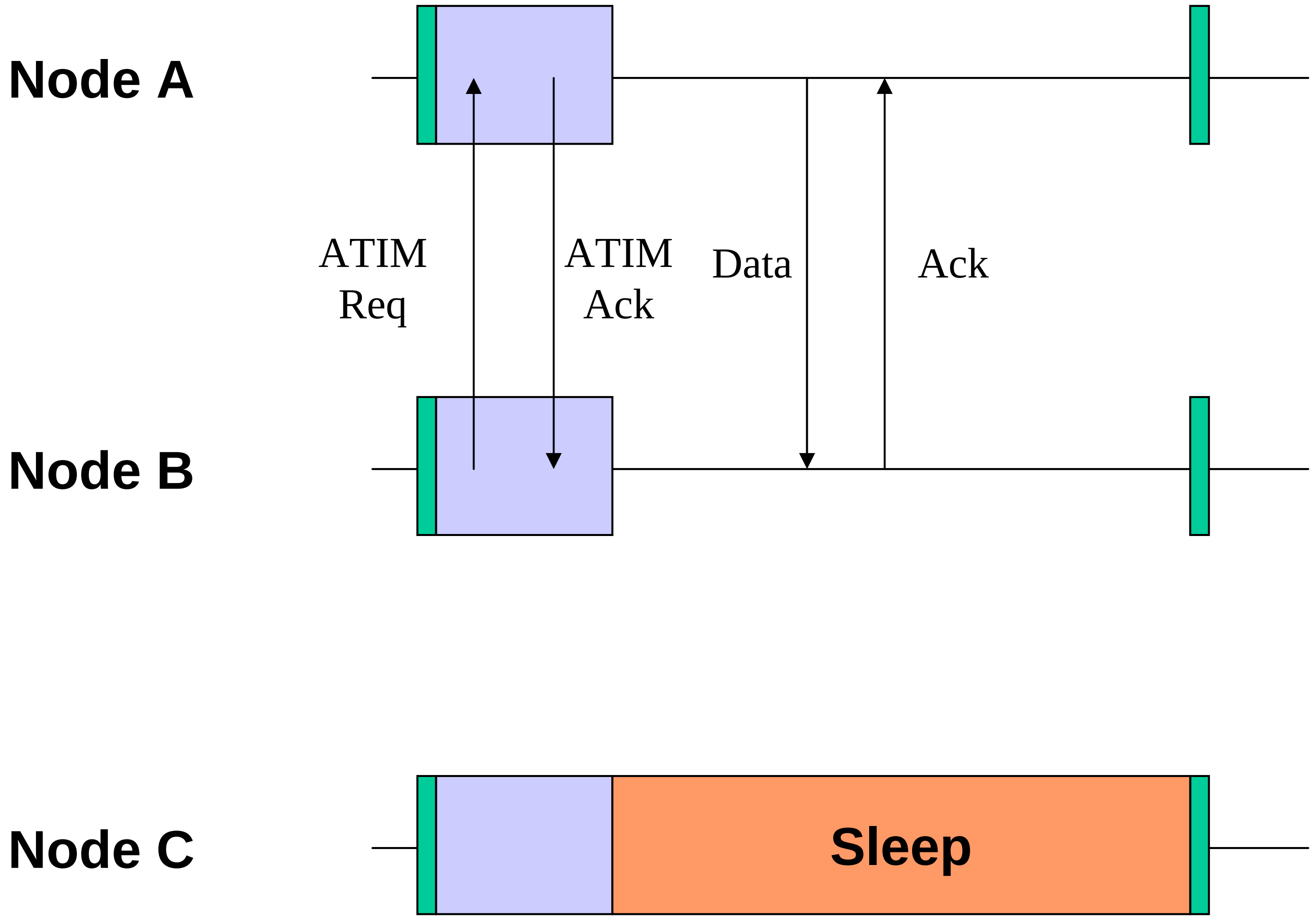

Power Save in IEEE 802.11 Ad Hoc Mode

- Size of ATIM window and beacon interval affects performance
- If ATIM window is too large, reduction in energy consumption reduced
  - Energy consumed during ATIM window
- If ATIM window is too small, not enough time to send ATIM request

Power Save in IEEE 802.11 Ad Hoc Mode

- How to choose ATIM window dynamically?
  - Based on observed load [Jung02infocom]
- How to synchronize hosts?
  - If two hosts’ ATIM windows do not overlap in time, they cannot exchange ATIM requests
  - Coordination requires that each host stay awake long enough (at least periodically) to discover out-of-sync neighbors [Tseng02infocom]

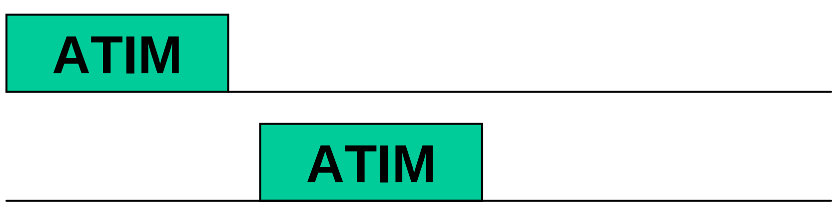

Impact on Upper Layers

- If each node uses the 802.11 power-save mechanism, each hop will require one beacon interval
  - This delay could be intolerable
- Allow upper layers to dictate whether a node should enter the power save mode or not [Chen01mobicom]

Energy Conservation

- Power save
- Power control

Power Control

- Power control has two potential benefits
  - Reduced interference & increased spatial reuse
  - Energy saving

Power Control

- When C transmits to D at a high power level, B cannot receive A's transmission due to interference from C

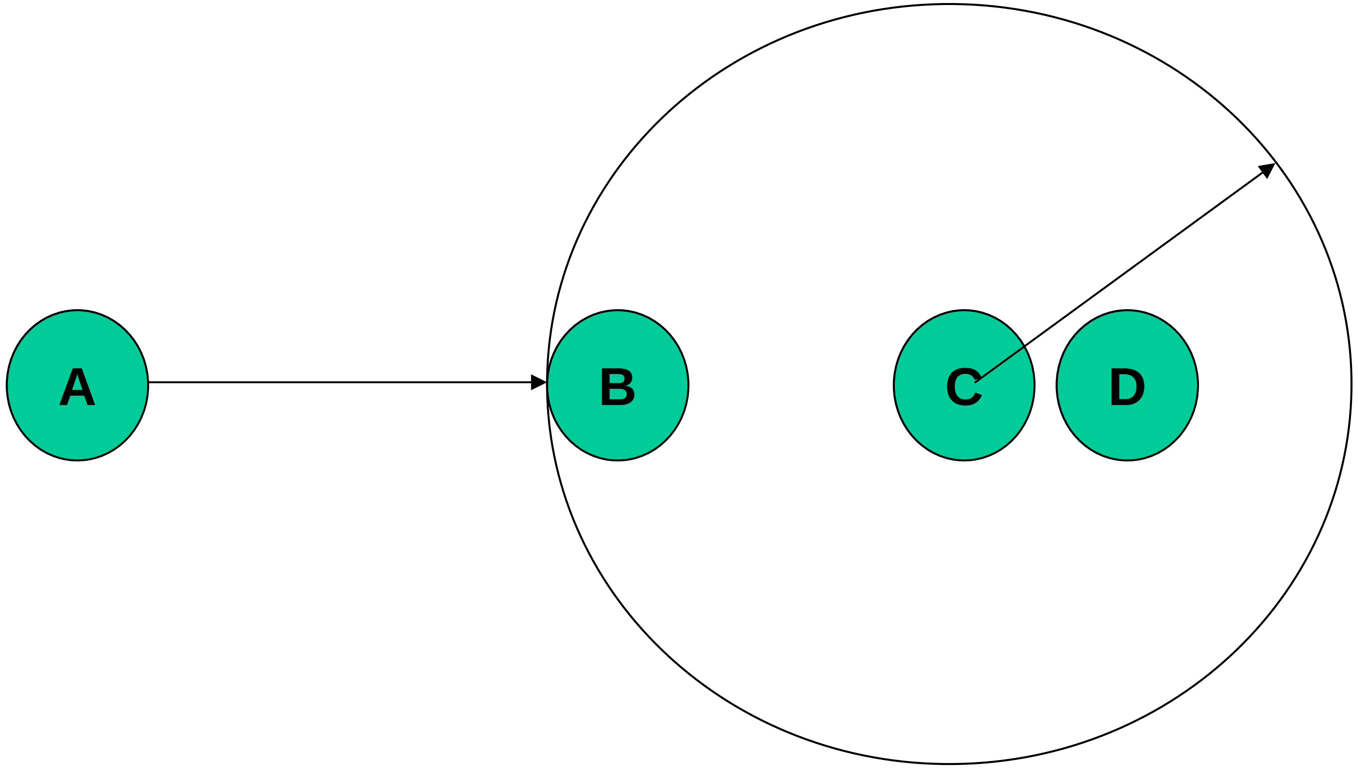

Power Control

- If C reduces transmit power, it can still communicate with D
  - Reduces energy consumption at node C
  - Allows B to receive A’s transmission (spatial reuse)

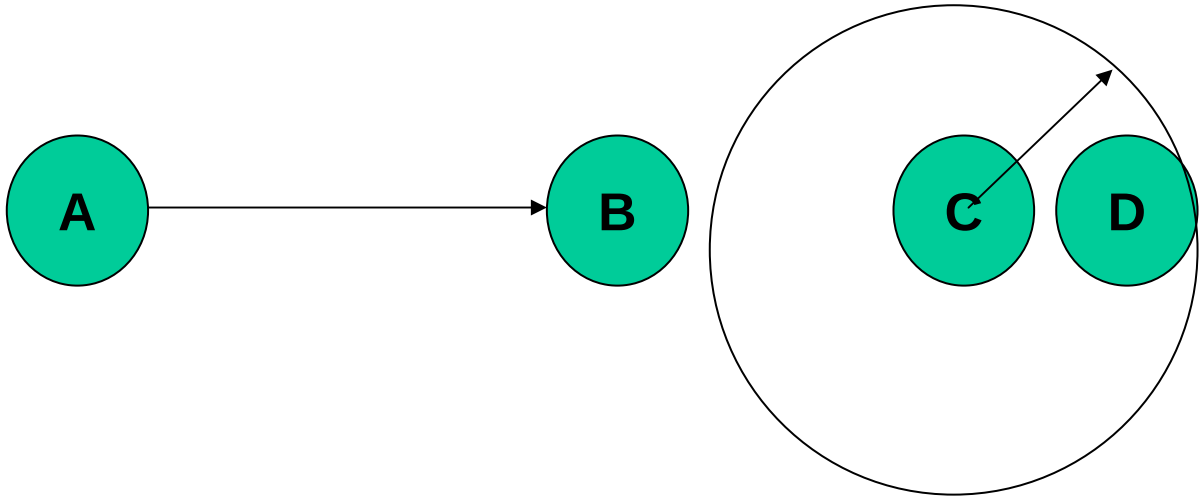

Power Control

- Received power level is proportional to $1/d^{\alpha}$, $\alpha >= 2$
- If power control is utilized, energy required to transmit to a host at distance d is proportional to
  d^{α} + constant
- Shorter hops typically preferred for energy consumption (depending on the constant) [Rodoplu99]
  - Transmit to C from A via B, instead of directly from A to C

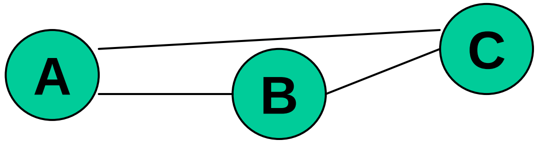

Power Control with 802.11

- Transmit RTS/CTS/DATA/ACK at least power level needed to communicate with the received

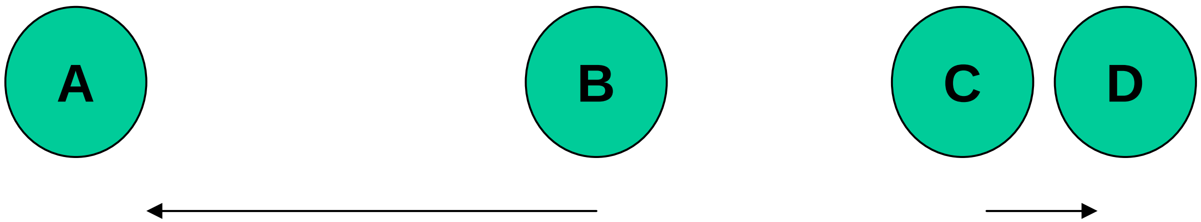

- A/B do not receive RTS/CTS from C/D. Also do not sense D's data transmission
- B's transmission to A at high power interferes with reception of ACK at C

A Plausible Solution

- RTS/CTS at highest power, and DATA/ACK at smallest necessary power level

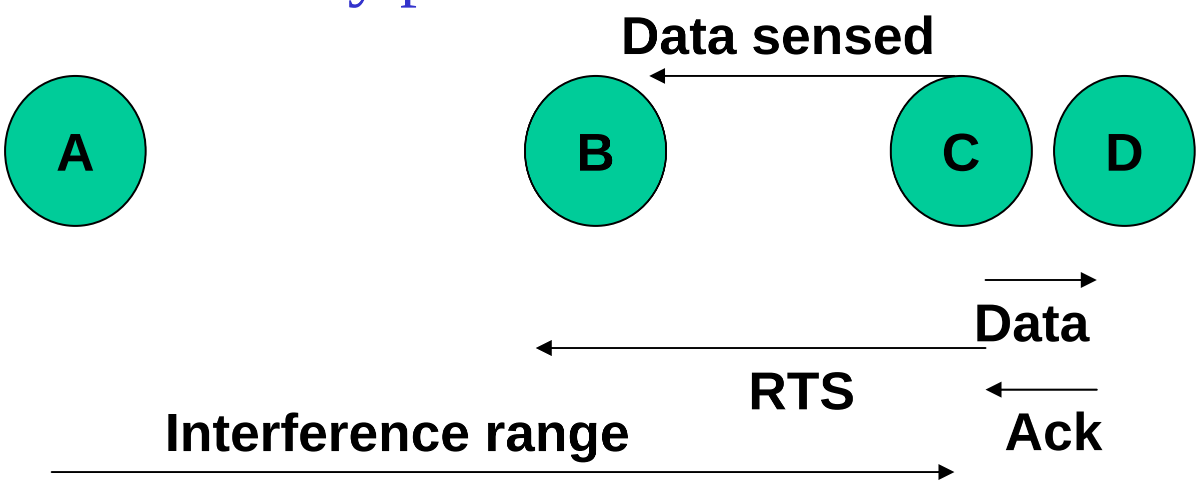

- A cannot sense C's data transmission, and may transmit DATA to some other host
- This DATA will interfere at C
- This situation unlikely if DATA transmitted at highest power level
  - Interference range ~ sensing range

Solution (cont.)

- Transmitting RTS at the highest power level also reduces spatial reuse
- Nodes receiving RTS/CTS have to defer transmissions

Caveat

- Energy saving by power control is limited to savings in transmit energy
- Other energy costs may not change, and may represent a significant fraction of total energy consumption

Power Controlled Multiple Access (PCMA) [Monks01infocom]

- If receiver node R can tolerate interference E, it sends a busy tone at power level C/E, where C is an appropriate constant
- When some node X receives a busy-tone a power level Pr, it may transmit at power level Pt <= C/Pr

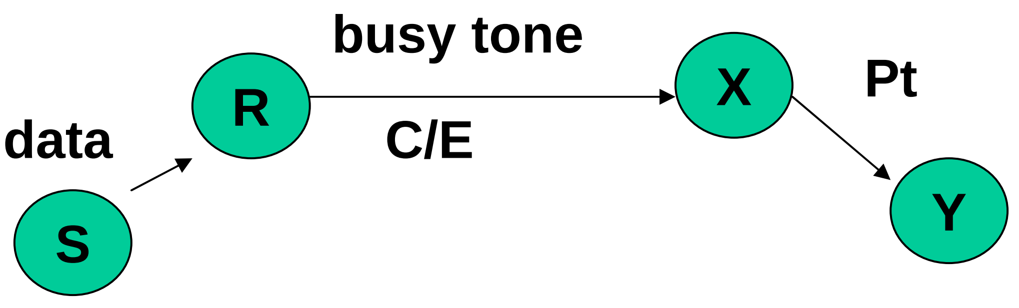

Power Controlled Multiple Access (PCMA)

- If receiver node R can tolerate noise E, it sends a busy tone at power level C/E, where C is an appropriate constant
- When some node X receives a busy-tone a power level Pr, it may transmit at power level Pt <= C/Pr
- Explanation:
  - Gain of channel RX = gain of channel XR = g
  - Busy tone signal level at X = Pr = g * C / E
  - Node X may transmit at level = Pt = C/Pr = E/g
  - Interference received by R = Pt * g = E

Power Controlled Multiple Access (PCMA)

- Advantage
  - Allows higher spatial reuse, as well as power saving using power control
- Disadvantages:
  - Need a separate channel for the busy tone
  - Since multiple nodes may transmit the busy tones simultaneously, spatial reuse is less than optimal

Adaptive Modulation

- Channel conditions are time-varying
- Received signal-to-noise ratio changes with time

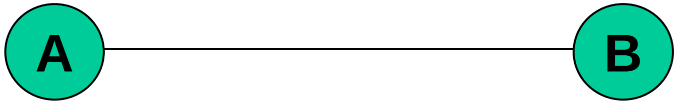

Adaptive Modulation

- Multi-rate radios are capable of transmitting at several rates, using different modulation schemes
- Choose modulation scheme as a function of channel conditions

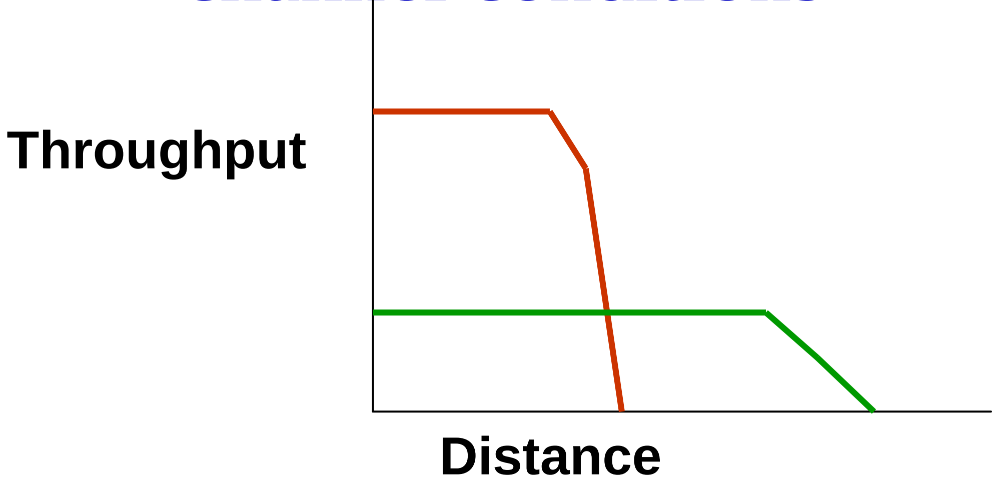

Modulation schemes provide a trade-off between throughput and range

Adaptive Modulation

- If physical layer chooses the modulation scheme transparent to MAC
  - MAC cannot know the time duration required for the transfer
- Must involve MAC protocol in deciding the modulation scheme
  - Some implementations use a sender-based scheme for this purpose [Kamerman97]
  - Receiver-based schemes can perform better

Sender-Based "Autorate Fallback" [Kamerman97]

- Probing mechanisms
- Sender decreases bit rate after X consecutive transmission attempts fail
- Sender increases bit rate after Y consecutive transmission attempt succeed

Autorate Fallback

- Advantage
  - Can be implemented at the sender, without making any changes to the 802.11 standard specification
- Disadvantage
  - Probing mechanism does not accurately detect channel state
  - Channel state detected more accurately at the receiver
  - Performance can suffer
    - Since the sender will periodically try to send at a rate higher than optimal
    - Also, when channel conditions improve, the rate is not increased immediately

Receiver-Based Autorate MAC [Holland01mobicom]

- Sender sends RTS containing its best rate estimate
- Receiver chooses best rate for the conditions and sends it in the CTS
- Sender transmits DATA packet at new rate
- Information in data packet header implicitly updates nodes that heard old rate

Receiver-Based Autorate MAC Protocol

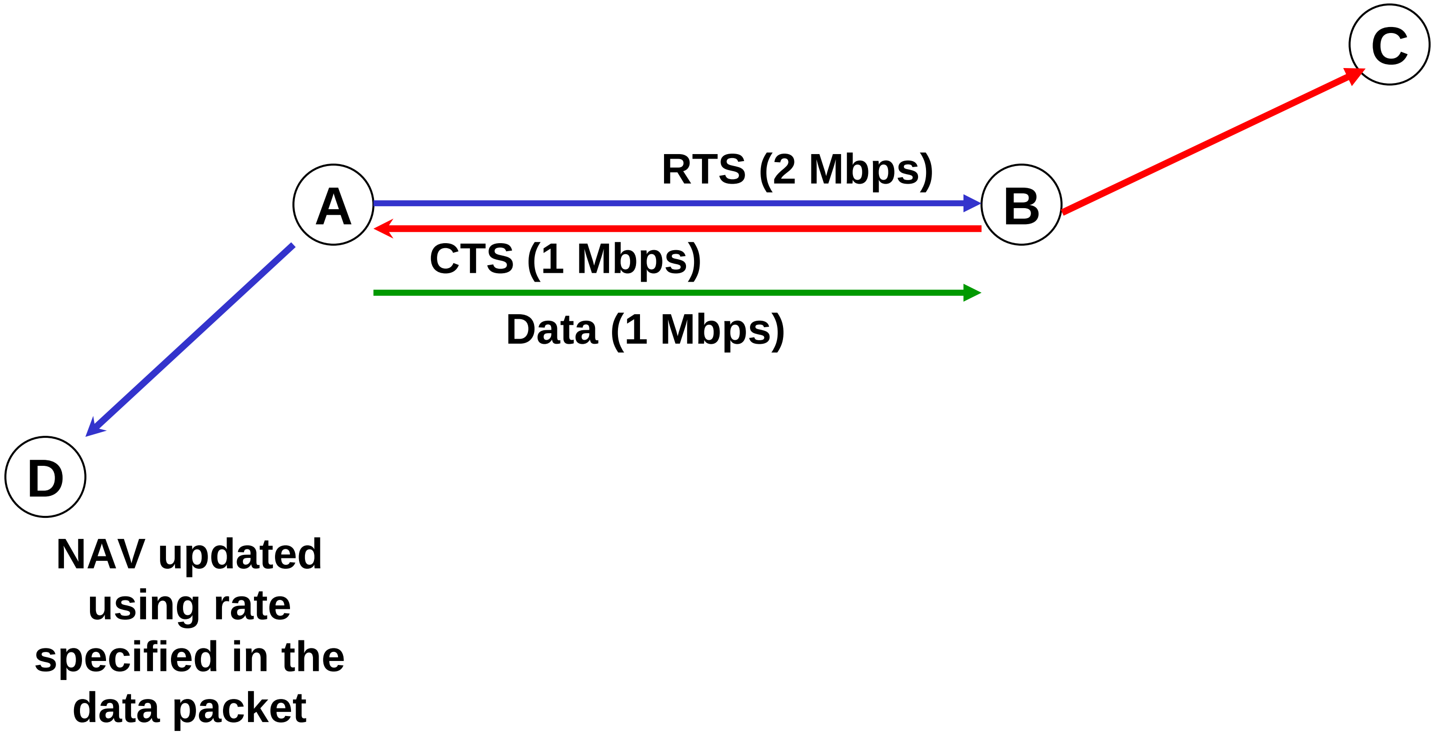

802.11b Physical Layer

Spectrum

- 802.11 operates in the unlicensed band (ISM – Industrial Scientific and Medical band) ~ 3 such bands
  - Cordless Telephony: 902 to 928 MHz
  - 802.11b: 2.4 to 2.483 GHz
  - 3rd ISM Band: 5.725 to 5.875 GHz
  - 802.11a: 5.15 to 5.825 GHz

Data Rates and Range

- 802.11: 2Mbps (Proposed in 1997)
- 802.11b: 1, 2, 5.5 and 11 Mbps, 100mts. range (product released in 1999, no product for 1 or 2 Mbps)
- 802.11g: 54Mbps, 100mts. range (uses OFDM)
- 802.11a: 6 to 54 Mbps, 50mts. range (uses OFDM)

802.11x

a → OFDM in the 5GHz band

b → High Rate DSSS in the 2.4GHz band

c → Bridge Operation Procedures

e → MAC Enhancements for QoS to improve QoS for better support of audio and video (such as MPEG-2) applications.

g → OFDM based 2.4 GHz WLAN.

i → Medium Access Method (MAC) Security Enhancements: enhance security and authentication mechanisms.

IEEE 802.11a

- 5 GHz (5.15-5.25, 5.25-5.35, 5.725-5.825 GHz) → ISM
- OFDM (Orthogonal Freq. Div. Multiplexing)
- 52 Subcarriers in OFDM
- BPSK/QPSK/QAM
- Forward Error Correction (Convolutional)
- Rates: 6, 9, 12, 18, 24, 36, 48, 54 Mbps

Base specifications

- Three Physical Layers:
  - FHSS (Frequency Hopping Spread Spectrum)
  - DSSS (Direct Sequence Spread Spectrum)
  - OFDM (Orthogonal Frequency Division Multiplexing)

Why Spread Spectrum?

- $C = B*\log_2(1+S/N)$
- To achieve the same channel capacity C
  - Large S/N, small B
  - Small S/N, large B
  - Increase S/N is inefficient due to the logarithmic relationship

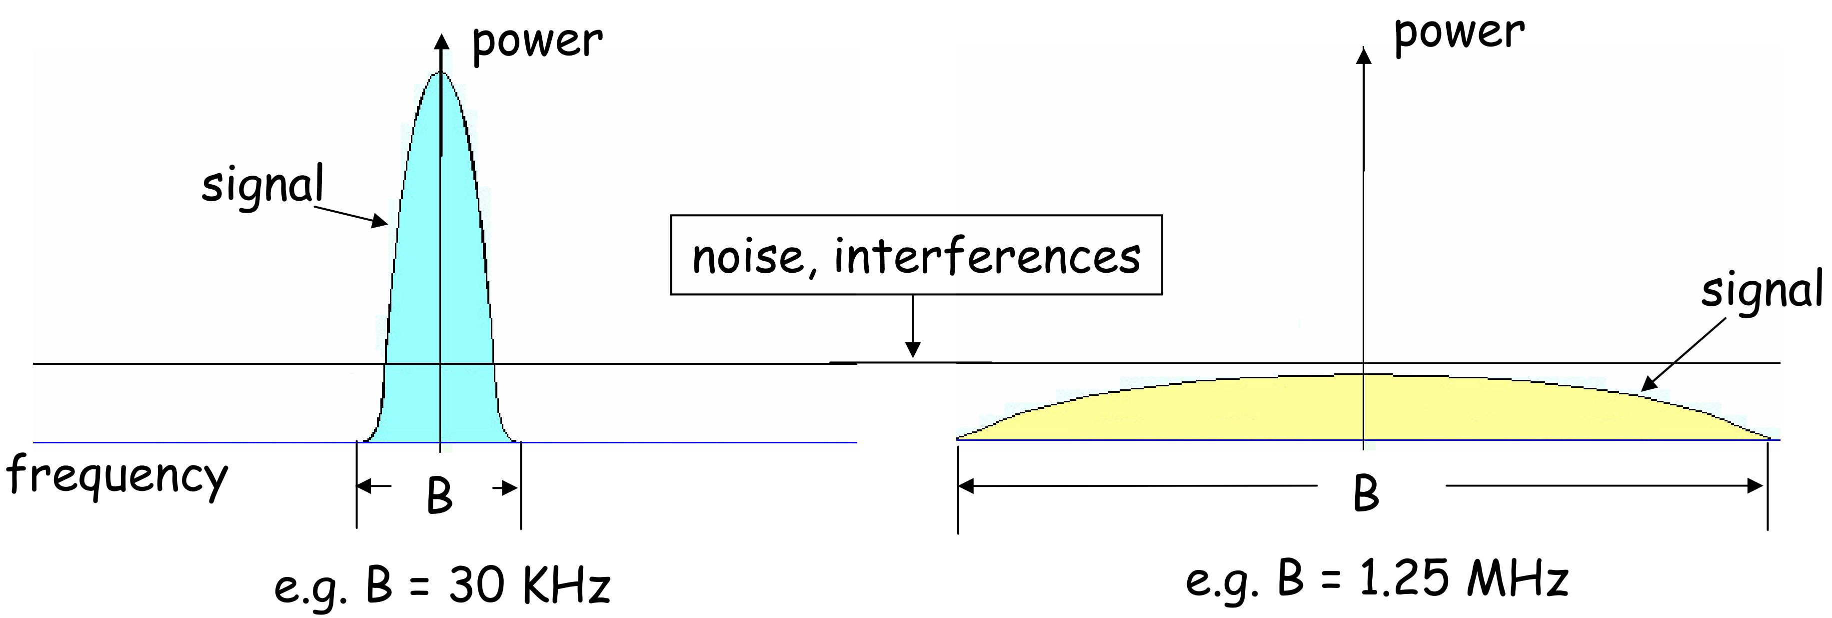

Frequency Hopping SS (FHSS)

- 2.4GHz band divided into 75 1MHz subchannels
- Sender and receive agree on a hopping pattern (pseudo random series). 22 hopping patterns defined

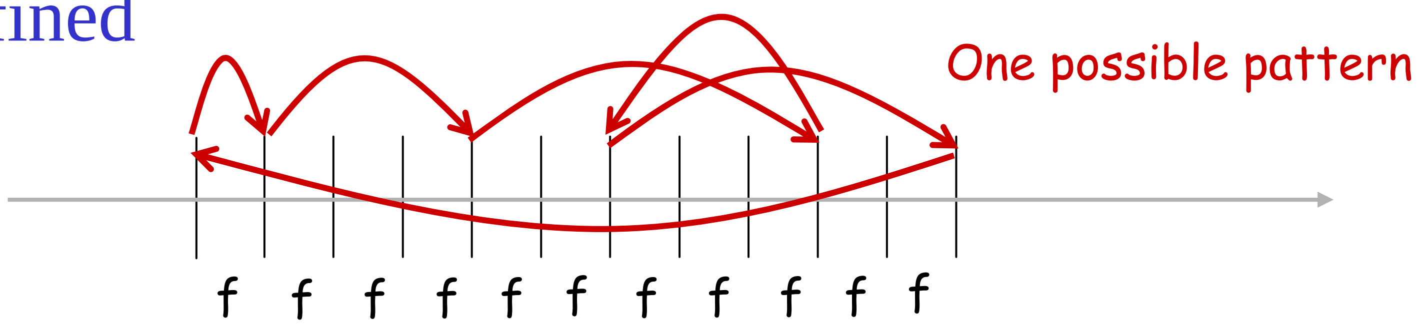

- Different hopping sequences enable co-existence of multiple BSSs
- Robust against narrow-band interferences

Direct Sequence SS

- Direct sequence (DS): most prevalent
  - Signal is spread by a wide bandwidth pseudorandom sequence (code sequence)
  - Signals appear as wideband noise to unintended receivers
- Not for intra-cell multiple access
  - Nodes in the same cell use same code sequence

802.11b PHY FRAME