CS/ECE 439: Wireless Networking

MAC Layer – Multi-Rate

 Fall 2020

What is "Data Rate" really?

- Number of bits that you transmit per unit time
  - under a fixed energy budget
- Too many bits/s
  - Each bit has little energy -> Hi BER
- Too few bits/s
  - Less BER but lower throughput

802.11b – Transmission rates

- Optimal rate depends on SINR
  - i.e., interference and current channel conditions

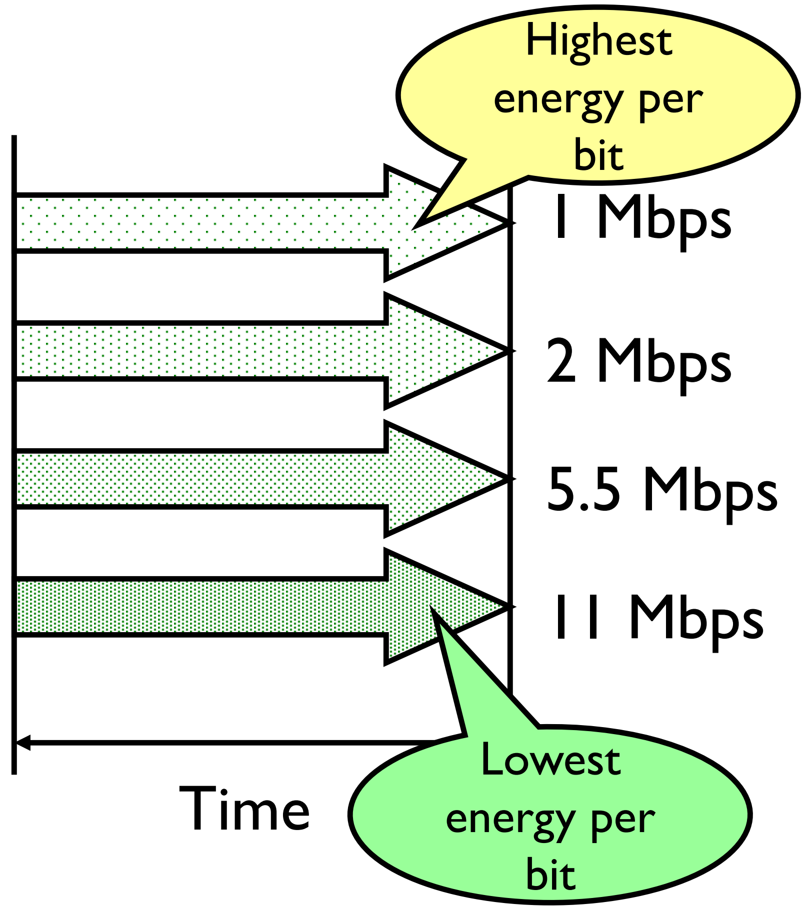

What is Multi-Rate?

- Ability of a wireless card to automatically operate at several different bit-rates
  - (e.g. 1, 2, 5.5, and 11 Mbps for 802.11b)
- Part of many existing wireless standards
  - (802.11b, 802.11a, 802.11g, HiperLAN2…)
- Virtually every wireless card in use today employs multi-rate

Example Carrier Modulations

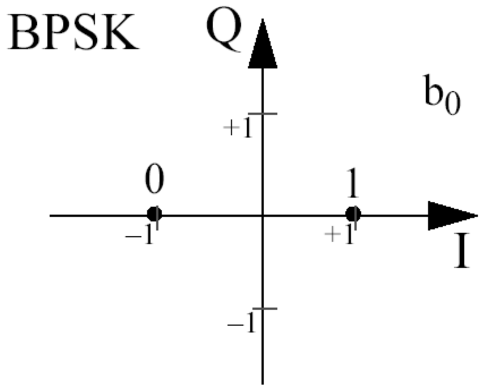

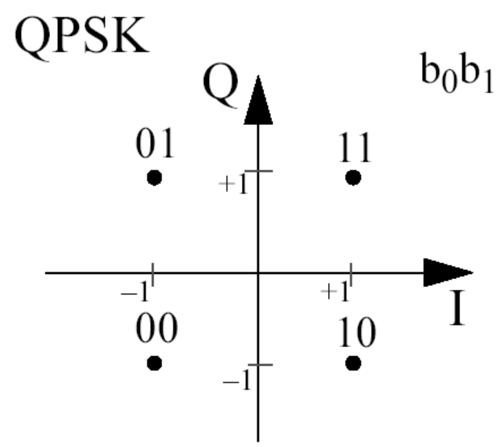

- Binary Phase Shift Keying
  - One bit per symbol
  - Made by the carrier and its inverse
- Quadrature Phase Shift Keying
  - Two bits per symbol
  - Uses quadrature carrier in addition to normal carrier
    - (90° phase shift of carrier)
  - 4 permutations for the inverse or not of the two carriers

Example Carrier Modulations (cont.)

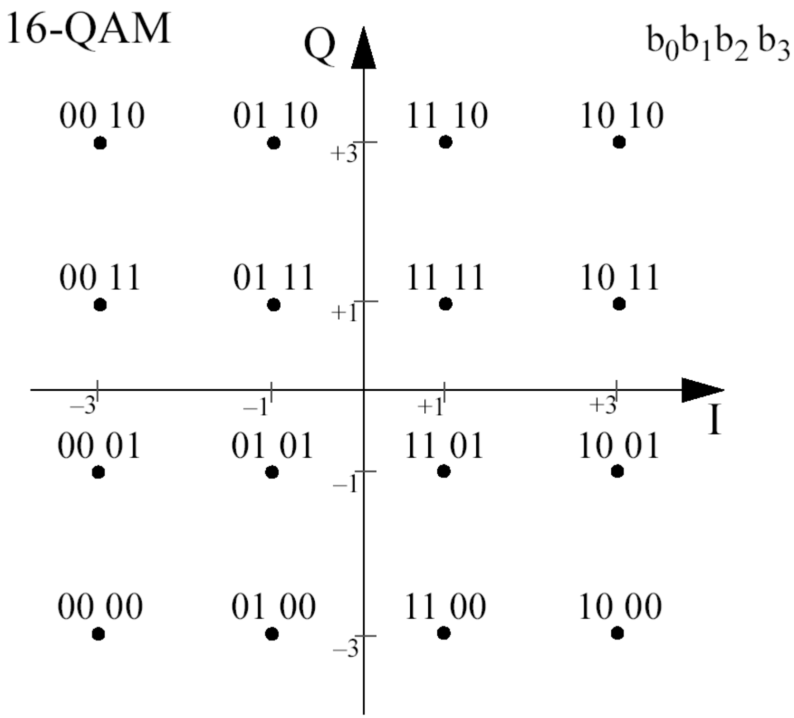

- 16 - Quadrature Amplitude Modulation
  - 4 bits per symbol
  - Also uses quadrature carrier
  - Each carrier is multiplied by +3, +1, -1, or -3
    - (amplitude modulation)
  - 16 possible combinations of the two multiplied carriers

Example Carrier Modulations (cont.)

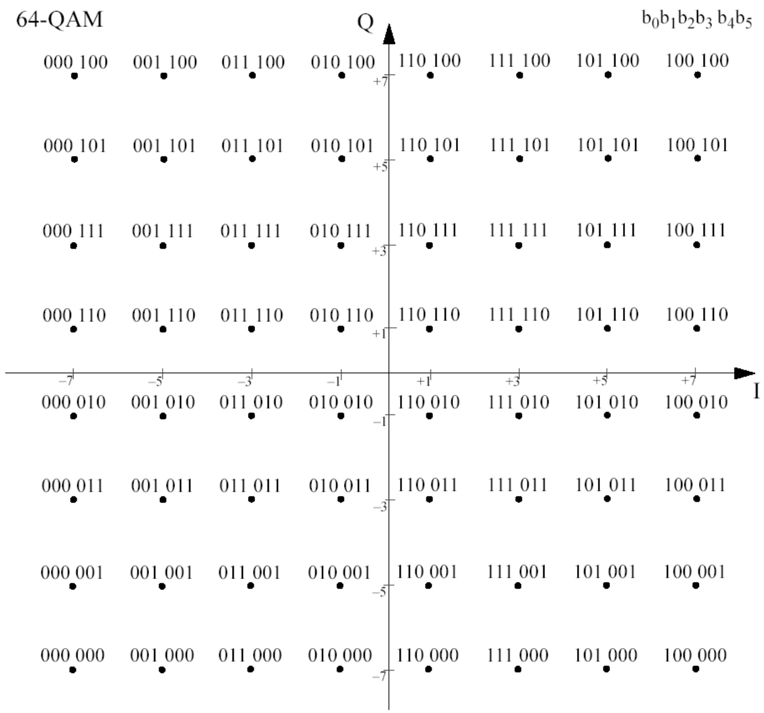

- **64 - Quadrature Amplitude Modulation**
  - 6 bits per symbol
  - Also uses quadrature carrier
  - Each carrier is multiplied by +7, +5, +3, +1, -1, -3, -5, or -7 (amplitude modulation)
  - 64 possible combinations of the two multiplied carriers

802.11a Rates resulting from Carrier Modulation and Coding

Data rate (Mbits/s)	Modulation	Coding rate (R)	Coded bits per subcarrier (N_{BPSC})	Coded bits per OFDM symbol (N_{CBPS})	Data bits per OFDM symbol (N_{DBPS})
6	BPSK	1/2	1	48	24
9	BPSK	3/4	1	48	36
12	QPSK	1/2	2	96	48
18	QPSK	3/4	2	96	72
24	16-QAM	1/2	4	192	96
36	16-QAM	3/4	4	192	144
48	64-QAM	2/3	6	288	192
54	64-QAM	3/4	6	288	216

Advantage of Multi-Rate?

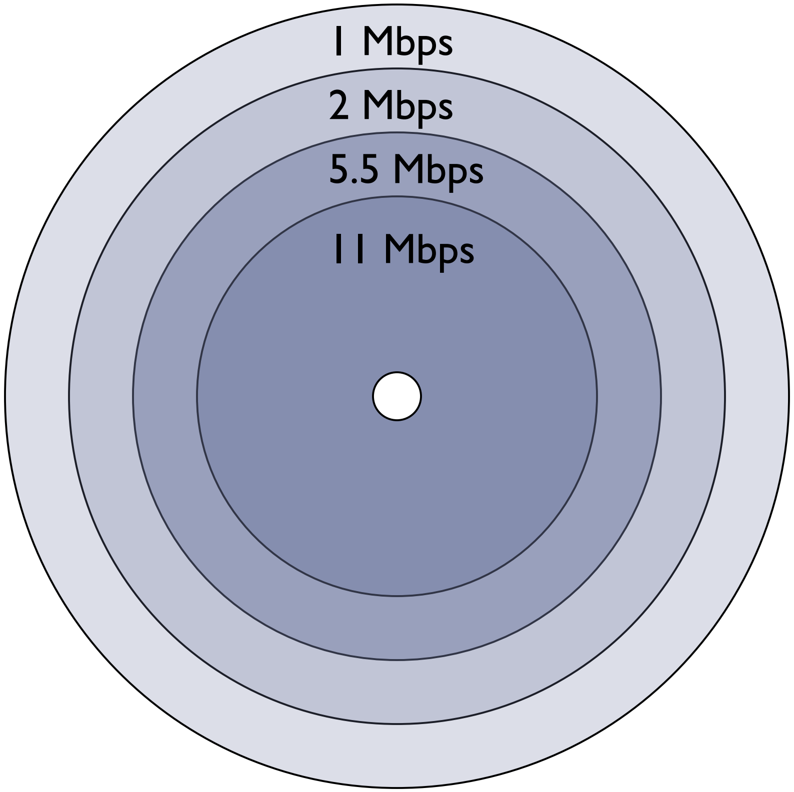

- Direct relationship between communication rate and the channel quality required for that rate
- As distance increases, channel quality decreases
  - Tradeoff between communication range and link speed
- Multi-rate provides flexibility to meet both consumer demands

Throughput vs. Distance for 802.11a

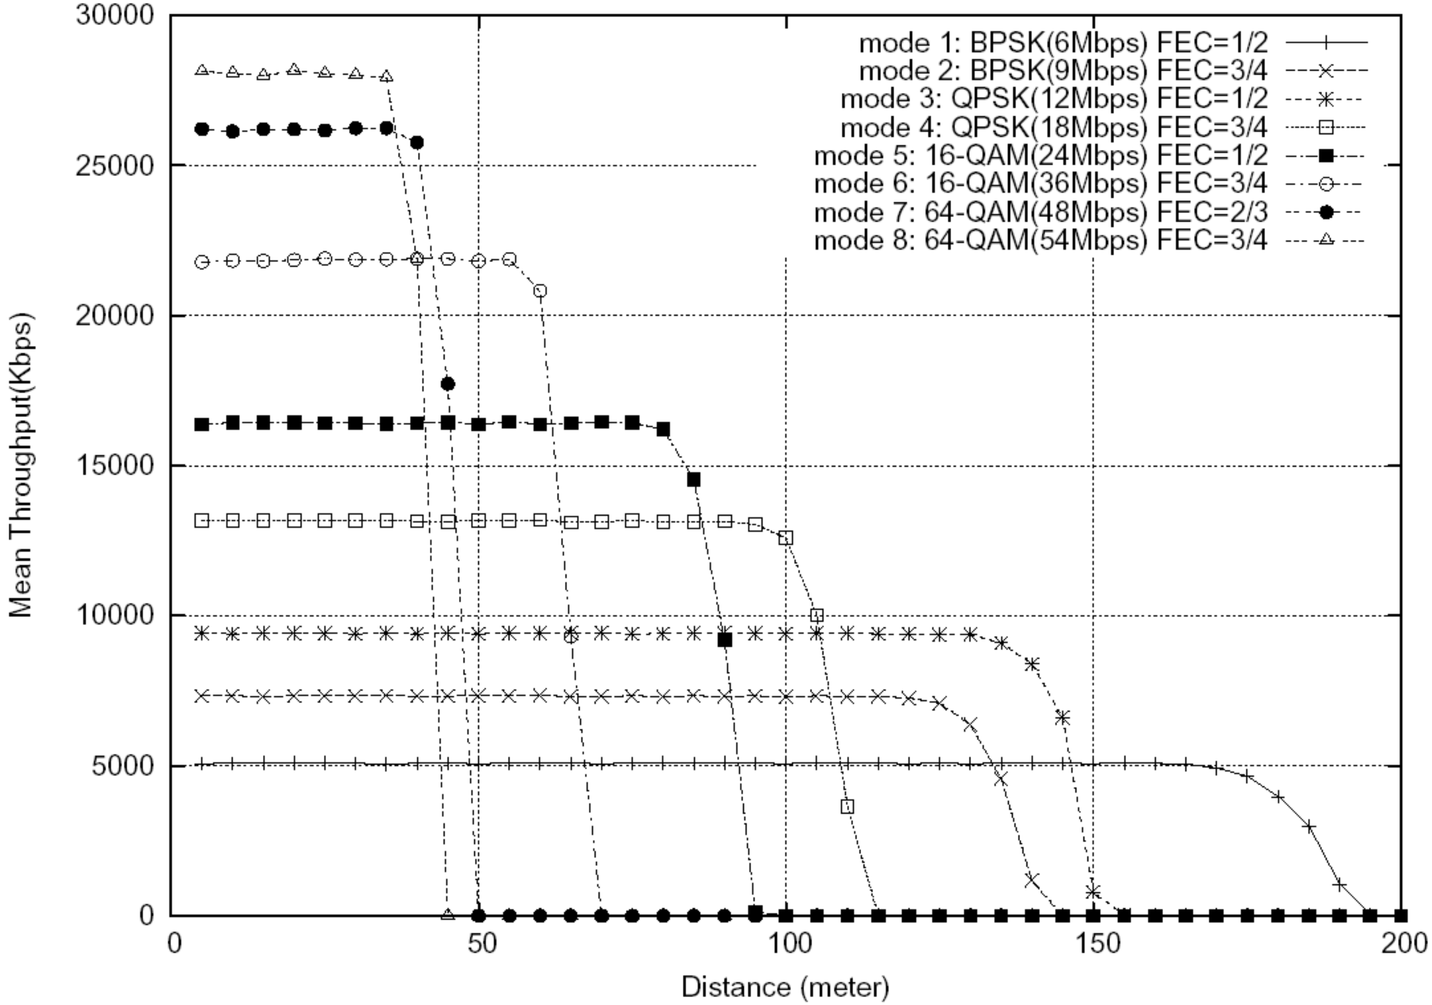

802.11 Frame Exchange Overhead

- Not all time is spent sending actual data

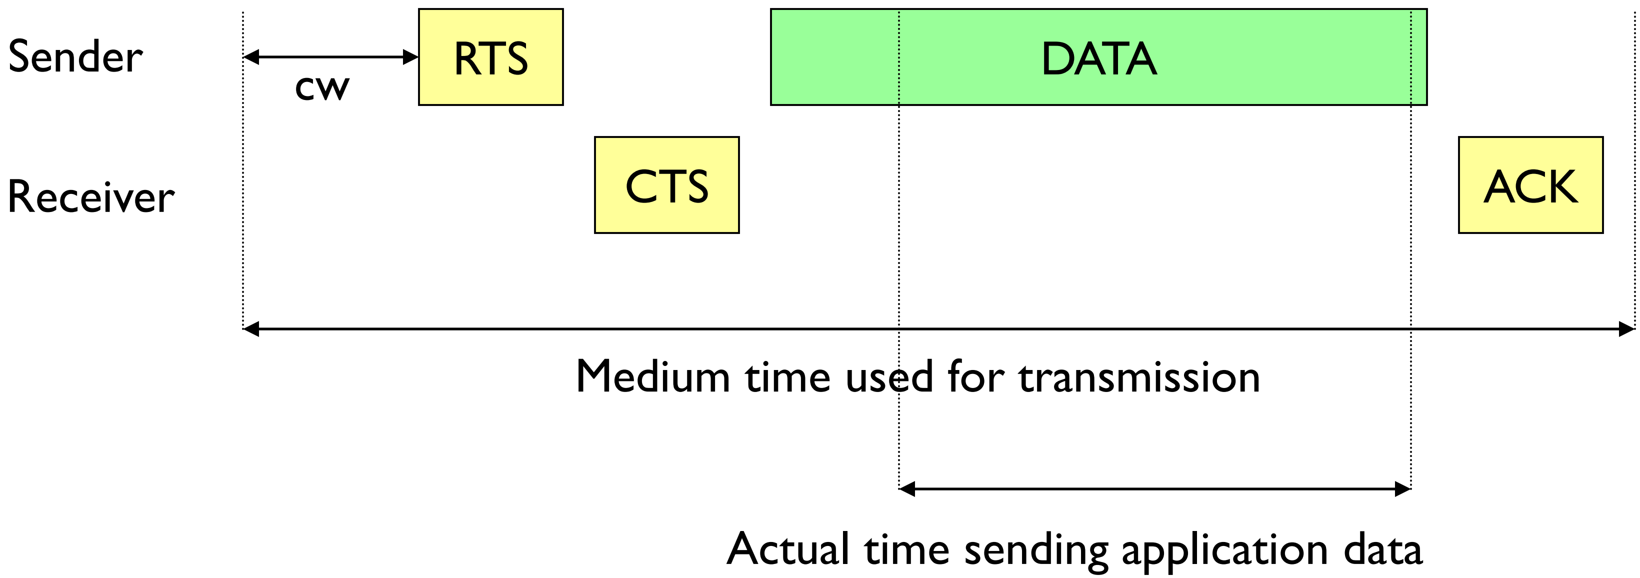

Multi-rate Frame in 802.11b

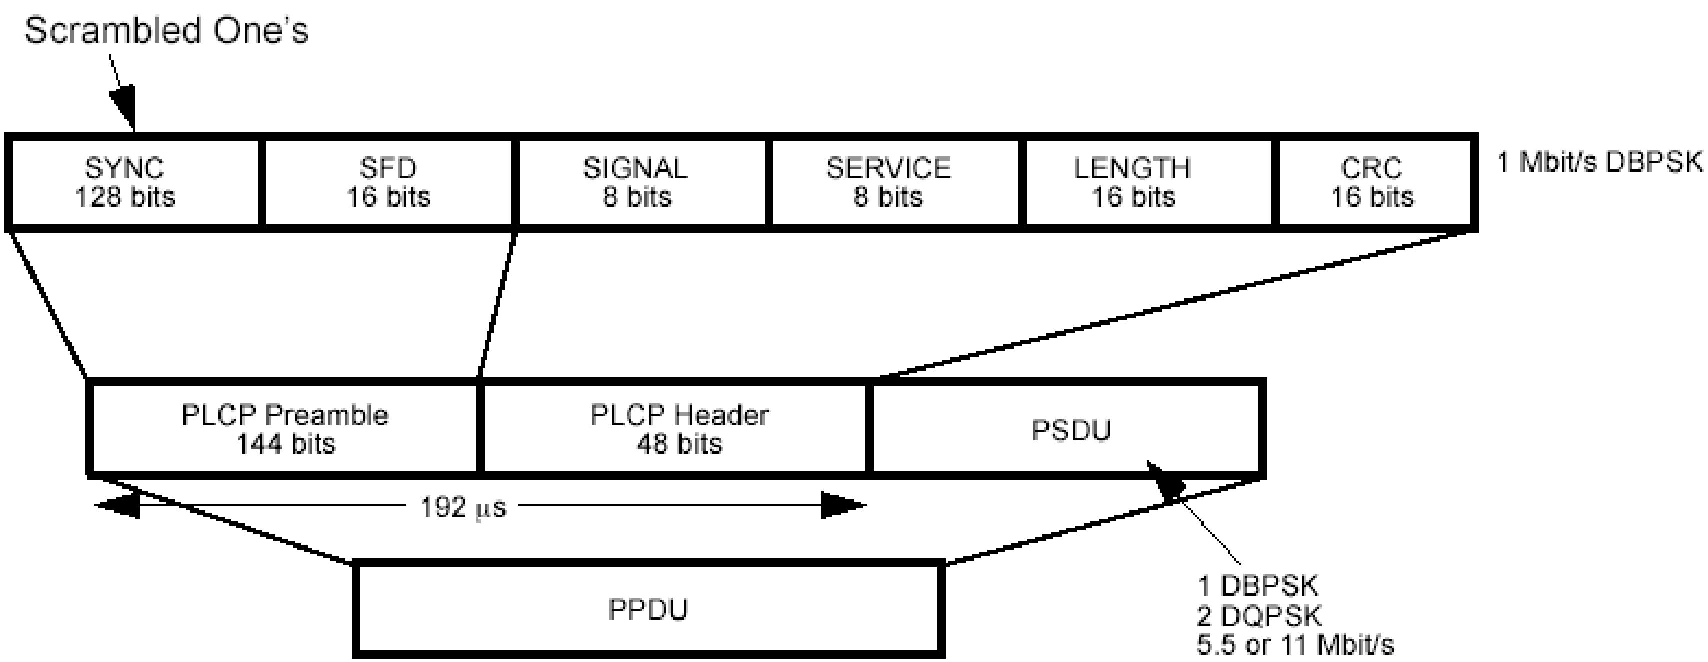

Figure 127—Long PLCP PPDU format

802.11b Frame Exchange Duration

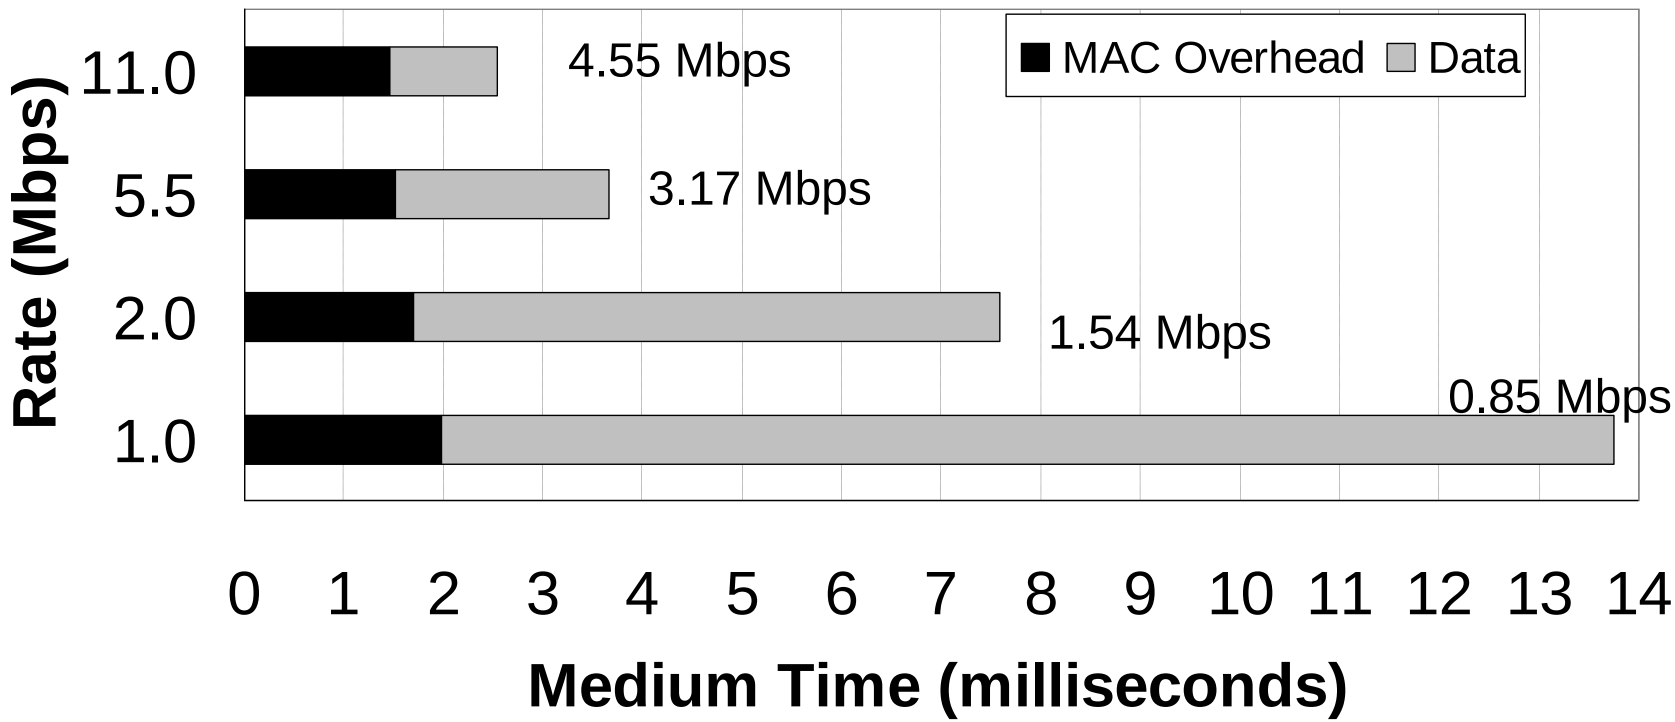

Medium Time consumed to transmit 1500 byte packet

Multi-rate Frame in 802.11a

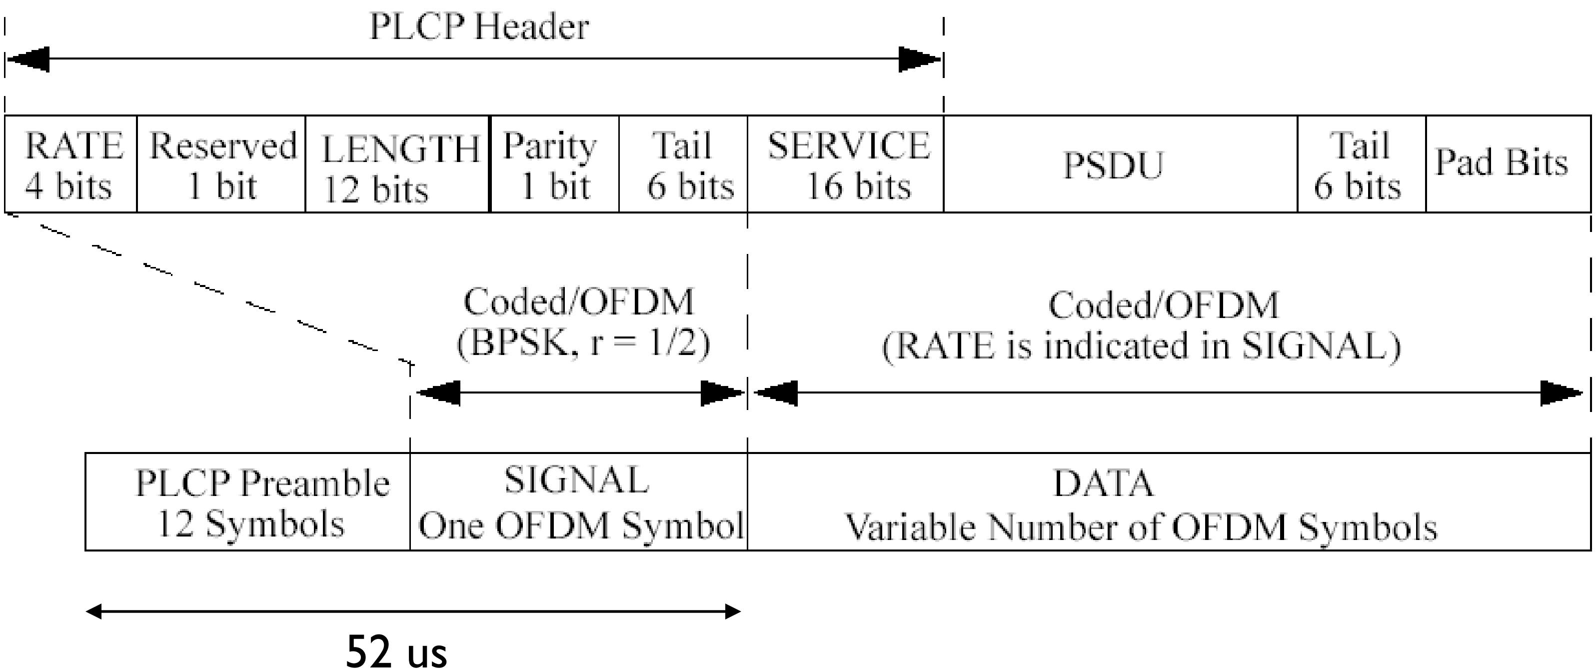

How do we choose modulation rates?

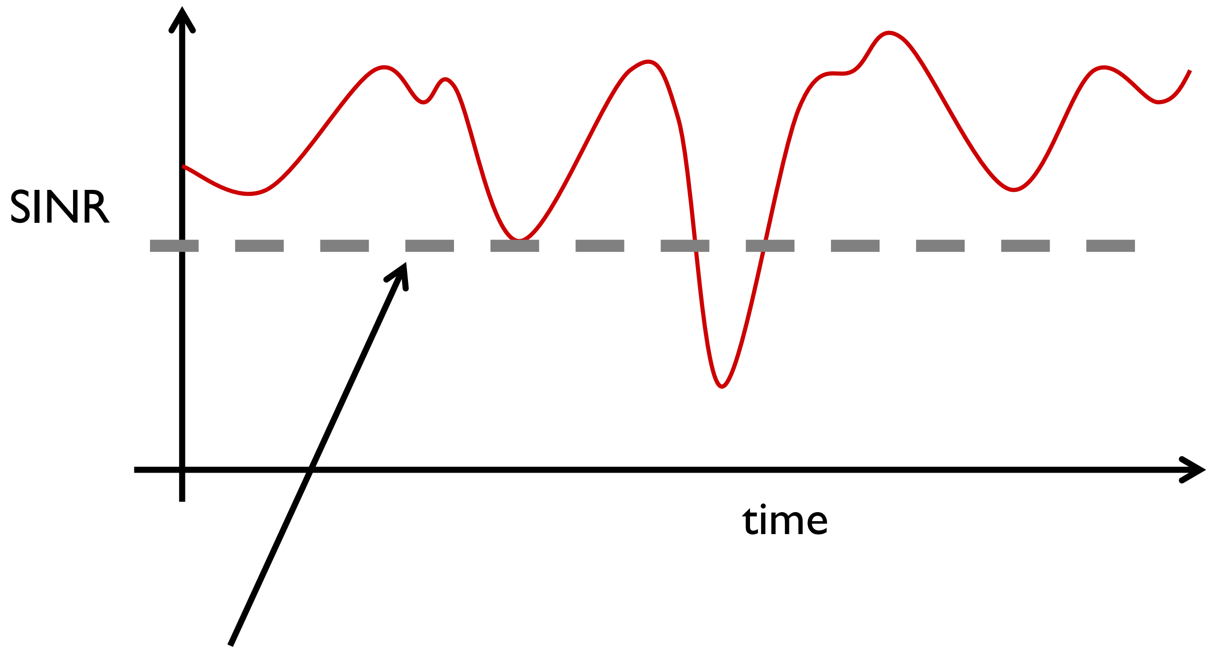

- Estimate a value of SINR
- Choose a corresponding rate that would transmit packets correctly most of the times
- Failure in some cases of fading
  - Live with it

Adaptive Rate-Control

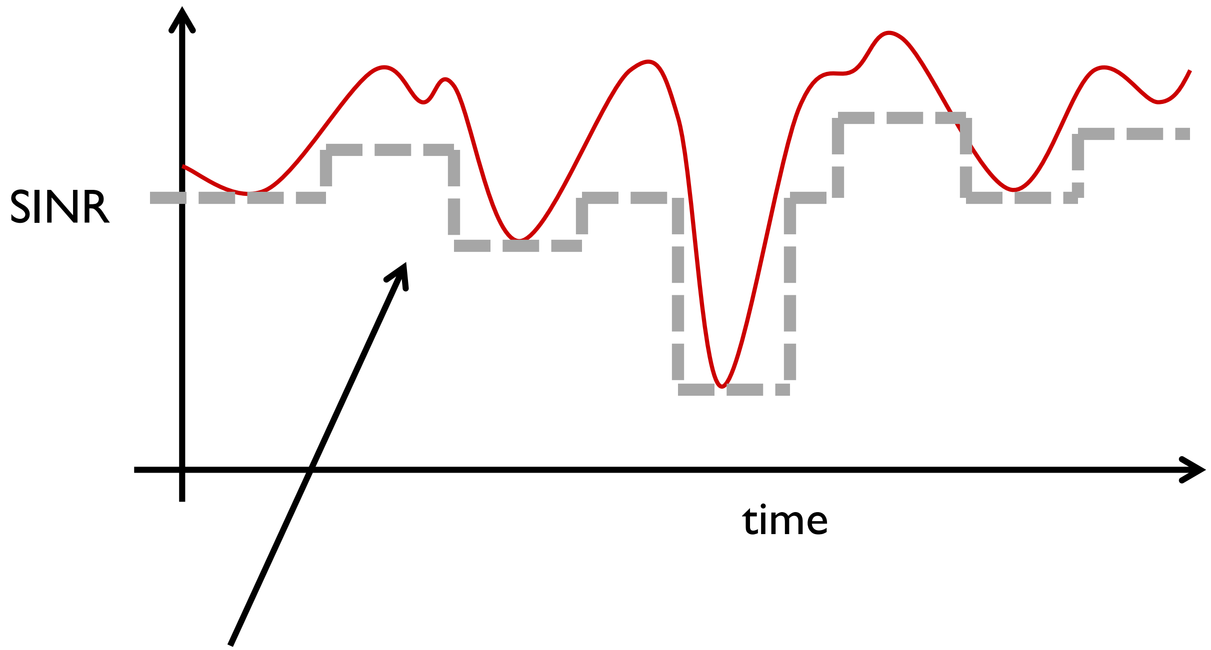

- Observe the current value of SINR
  - Use as indicator of near-future value
- Choose corresponding rate of modulation
- Repeat
  - Controls rate if channel conditions have changed

Seems simple, but …

- Rate control has variety of implications
  - Any single MAC protocol solves part of the puzzle
- Important to understand e2e implications
  - Does routing protocols get affected?
  - Does TCP get affected?
  - …
- Good to make a start at the MAC layer
  - ARF
  - RBAR
  - OAR
  - …

Problem

- Modulation schemes have different error characteristics

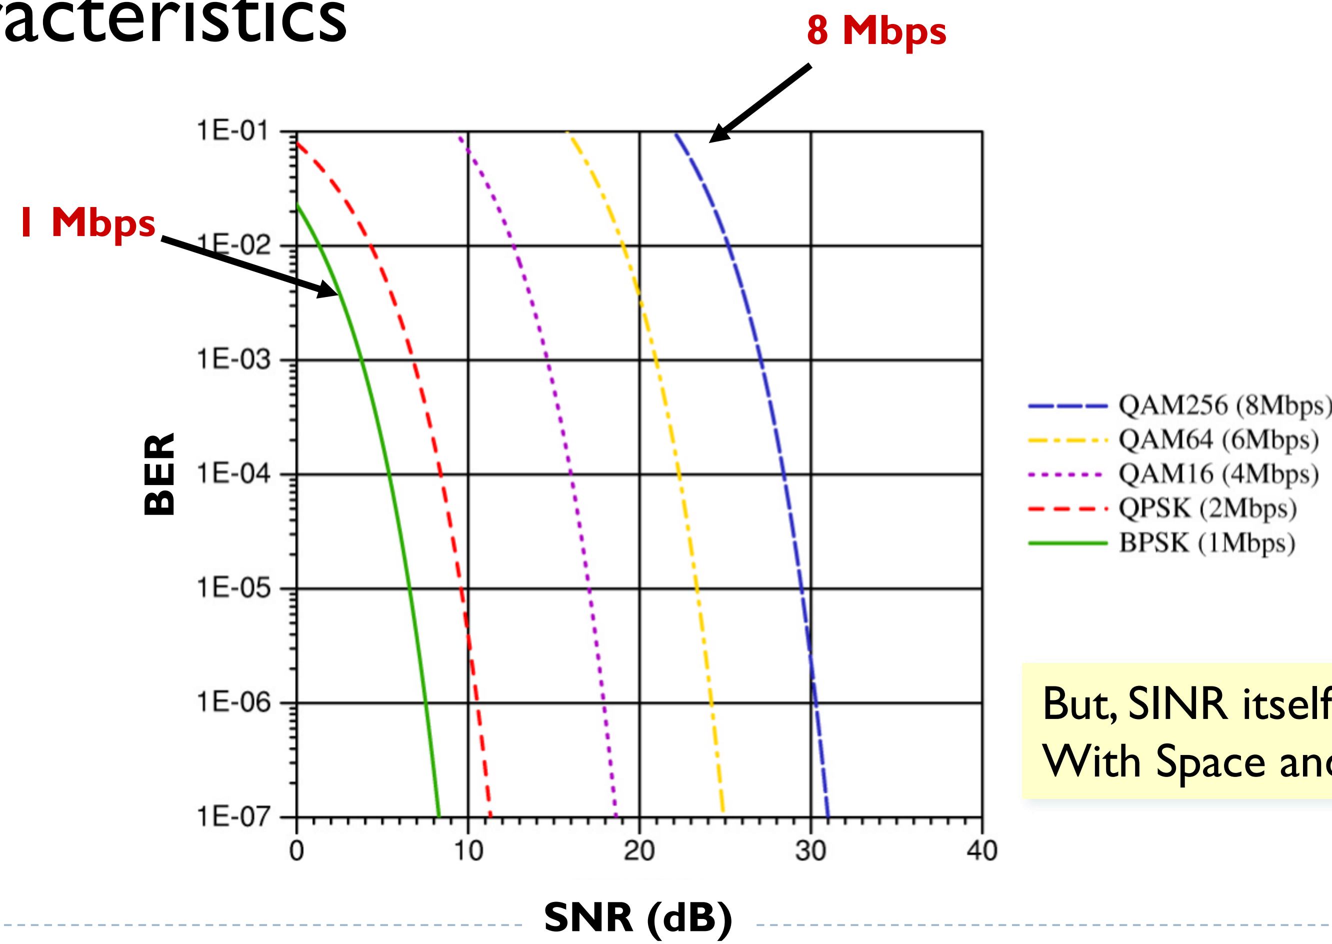

But, SINR itself varies With Space and Time

Impact

- Large-scale variation with distance (Path loss)

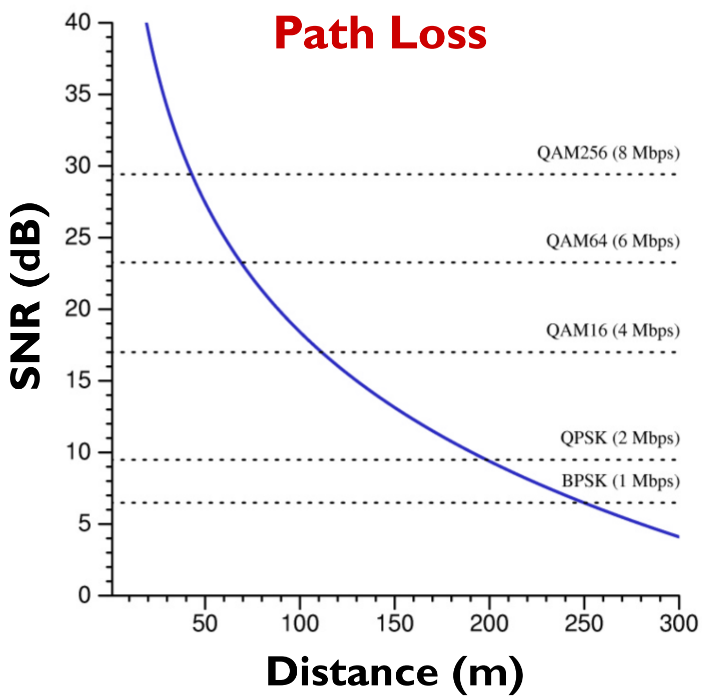

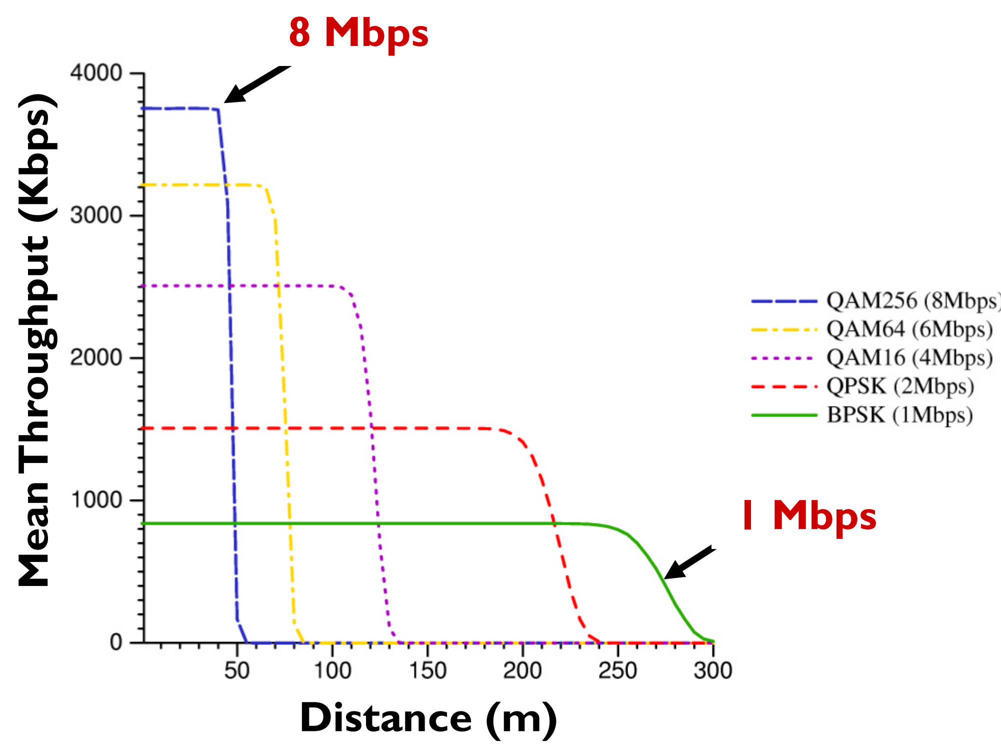

Impact

- Small-scale variation with time (Fading)

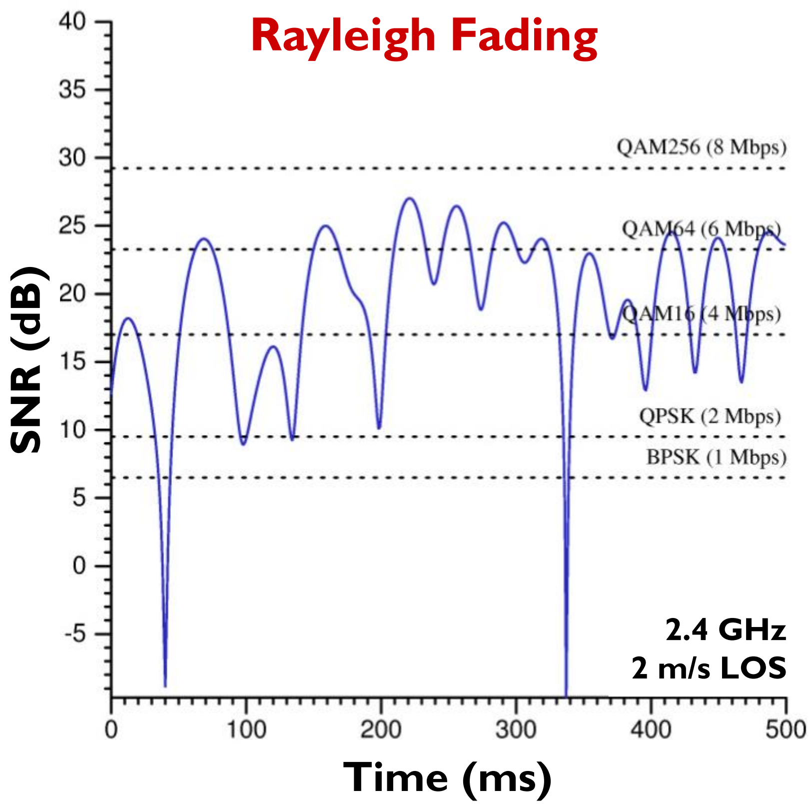

Which modulation scheme is best?

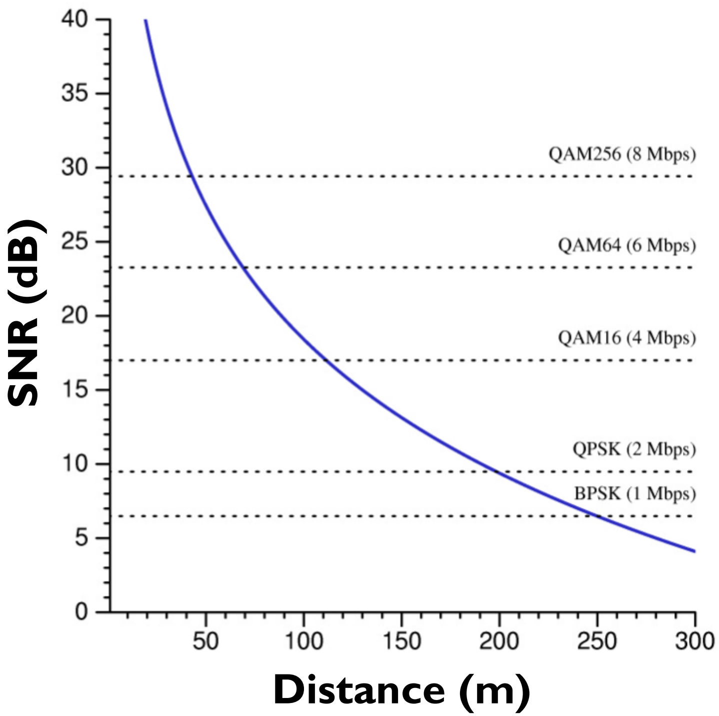

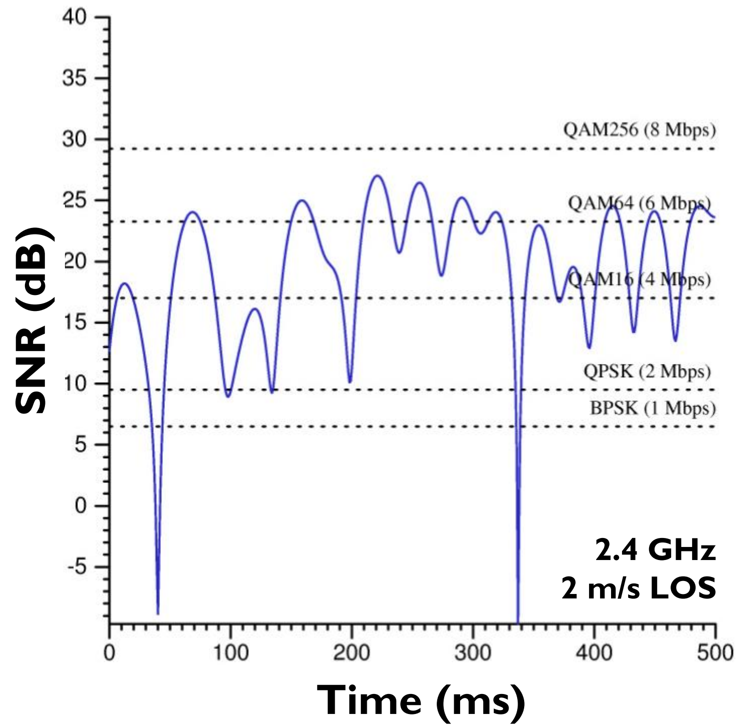

Answer → Rate Adaptation

- Dynamically choose the best modulation scheme for the channel conditions

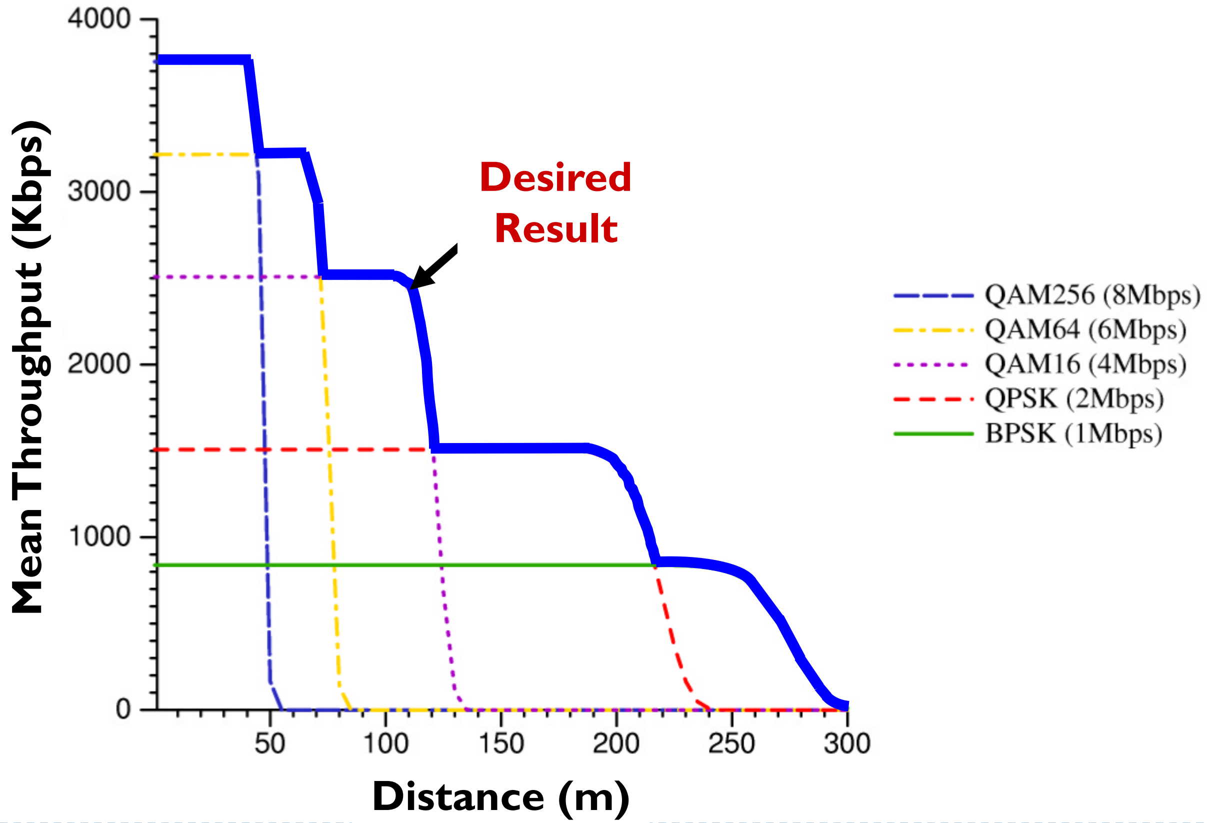

Design Issues

- How frequently should we adapt the rate?
  - Signal can vary rapidly depending on
    - carrier frequency
    - node speed
    - interference
    - etc.
- For conventional hardware at pedestrian speeds, rate adaptation is feasible on a per-packet basis

Adaptation → At Which Layer ?

- Cellular networks
  - Adaptation at the physical layer
- Impractical for 802.11 in WLANs ← Why?

Adaptation → At Which Layer ?

- Cellular networks
  - Adaptation at the physical layer
- Impractical for 802.11 in WLANs ← Why?

RTS/CTS requires that the *rate be known in advance*

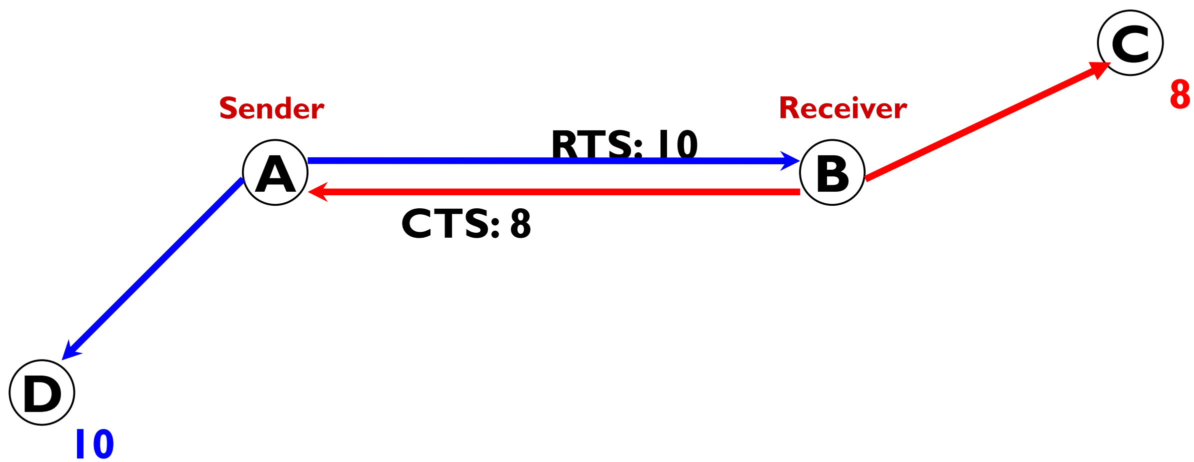

- For WLANs, rate adaptation is best handled at the MAC layer

Who should select the data rate?

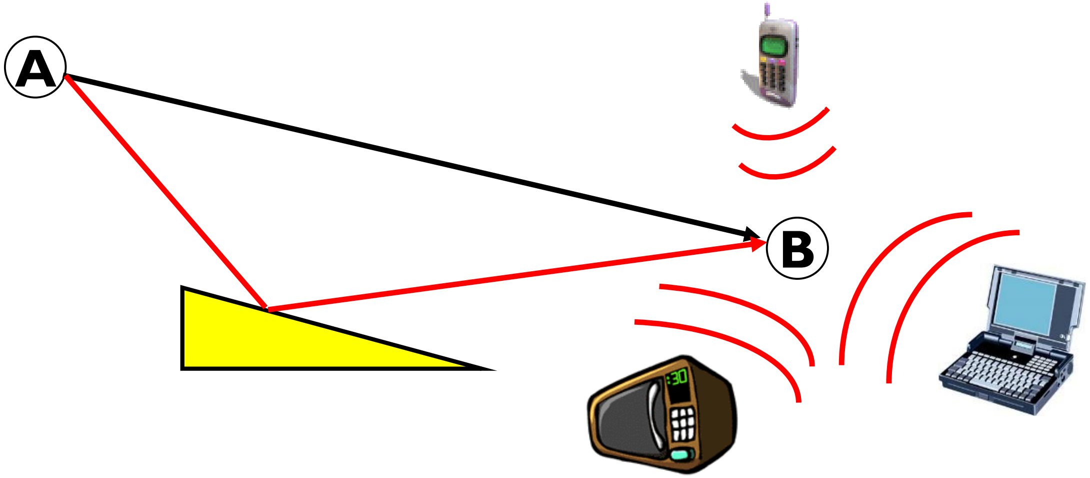

Who should select the data rate?

- Collision is at the receiver
- Channel conditions are only known at the receiver
  - SS, interference, noise, BER, etc.

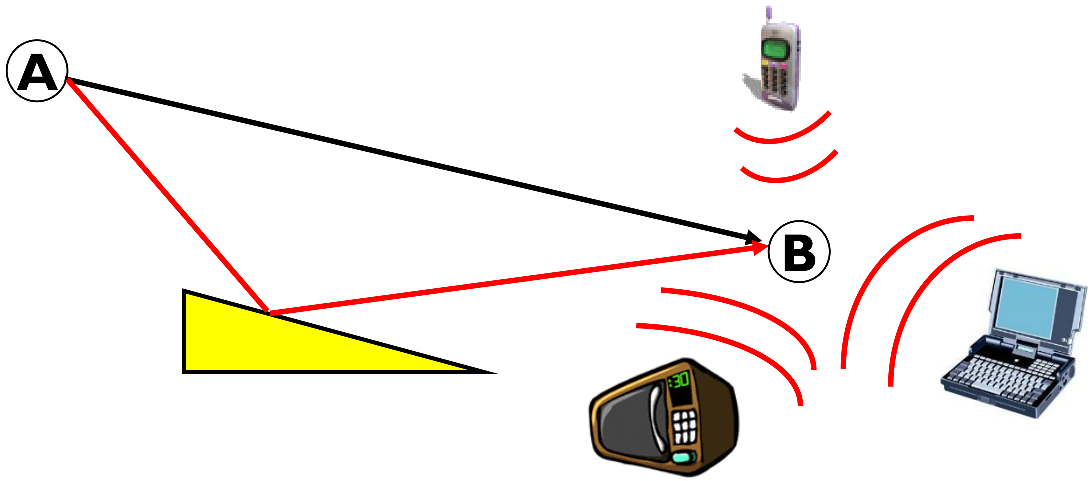

- The receiver is best positioned to select data rate

Lucent WaveLAN "Autorate Fallback" (ARF)

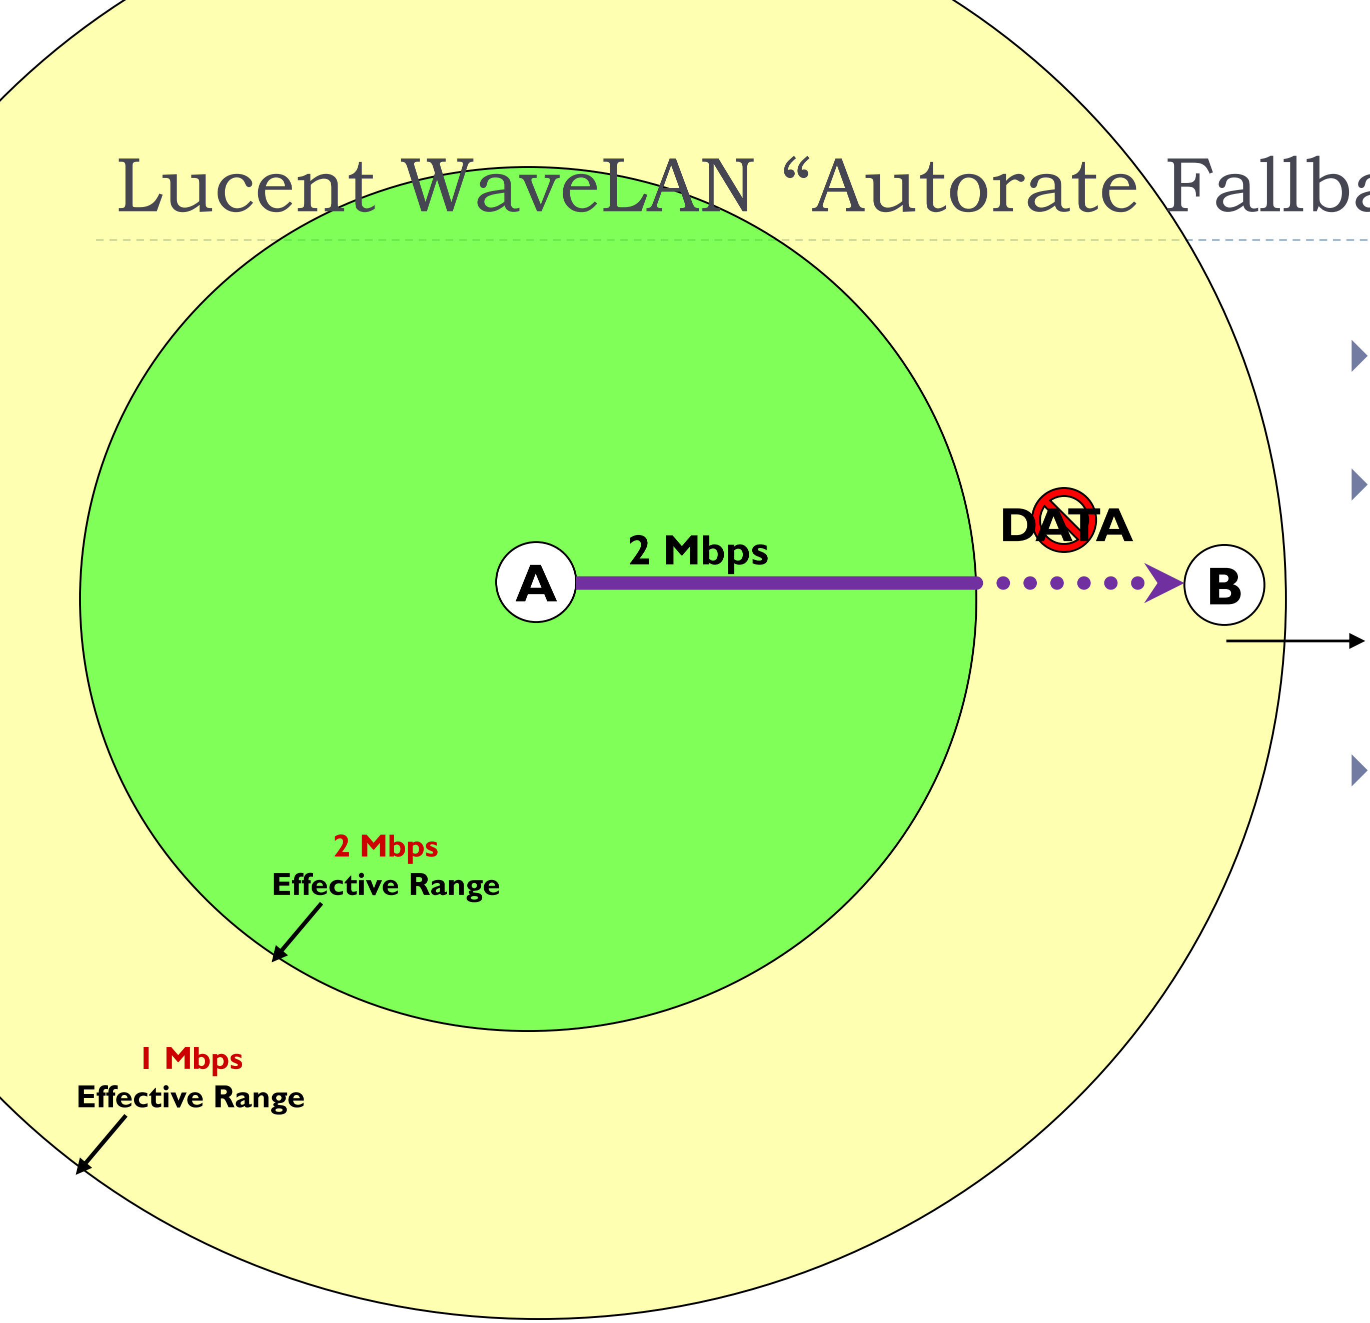

- Lost ACKs indicate link quality
- Sender decreases rate after
  - N consecutive ACKS are lost
- Sender increases rate after
  - Y consecutive ACKS are received or
  - T secs have elapsed since last attempt

Performance of ARF

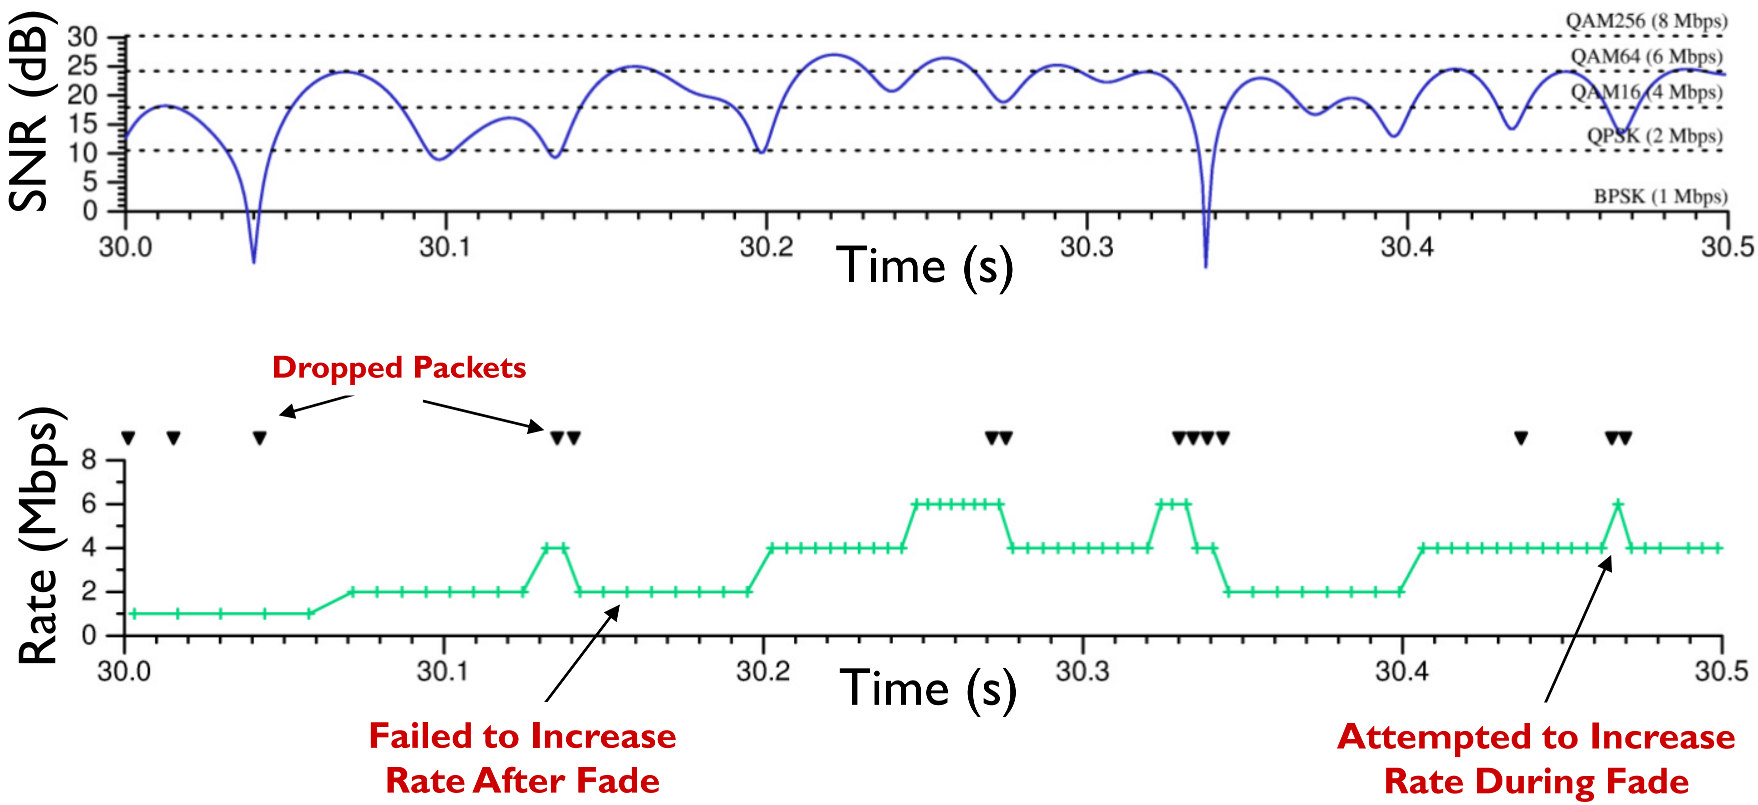

- Slow to adapt to channel conditions
- Choice of N, Y, T may not be best for all situations

Receiver-Based Autorate (RBAR)

- Move the rate adaptation mechanism to the receiver
  - Better channel quality information = better rate selection
- Utilize the RTS/CTS exchange to
  - Provide the receiver with a signal to sample (RTS)
  - Carry feedback (data rate) to the sender (CTS)

RBAR

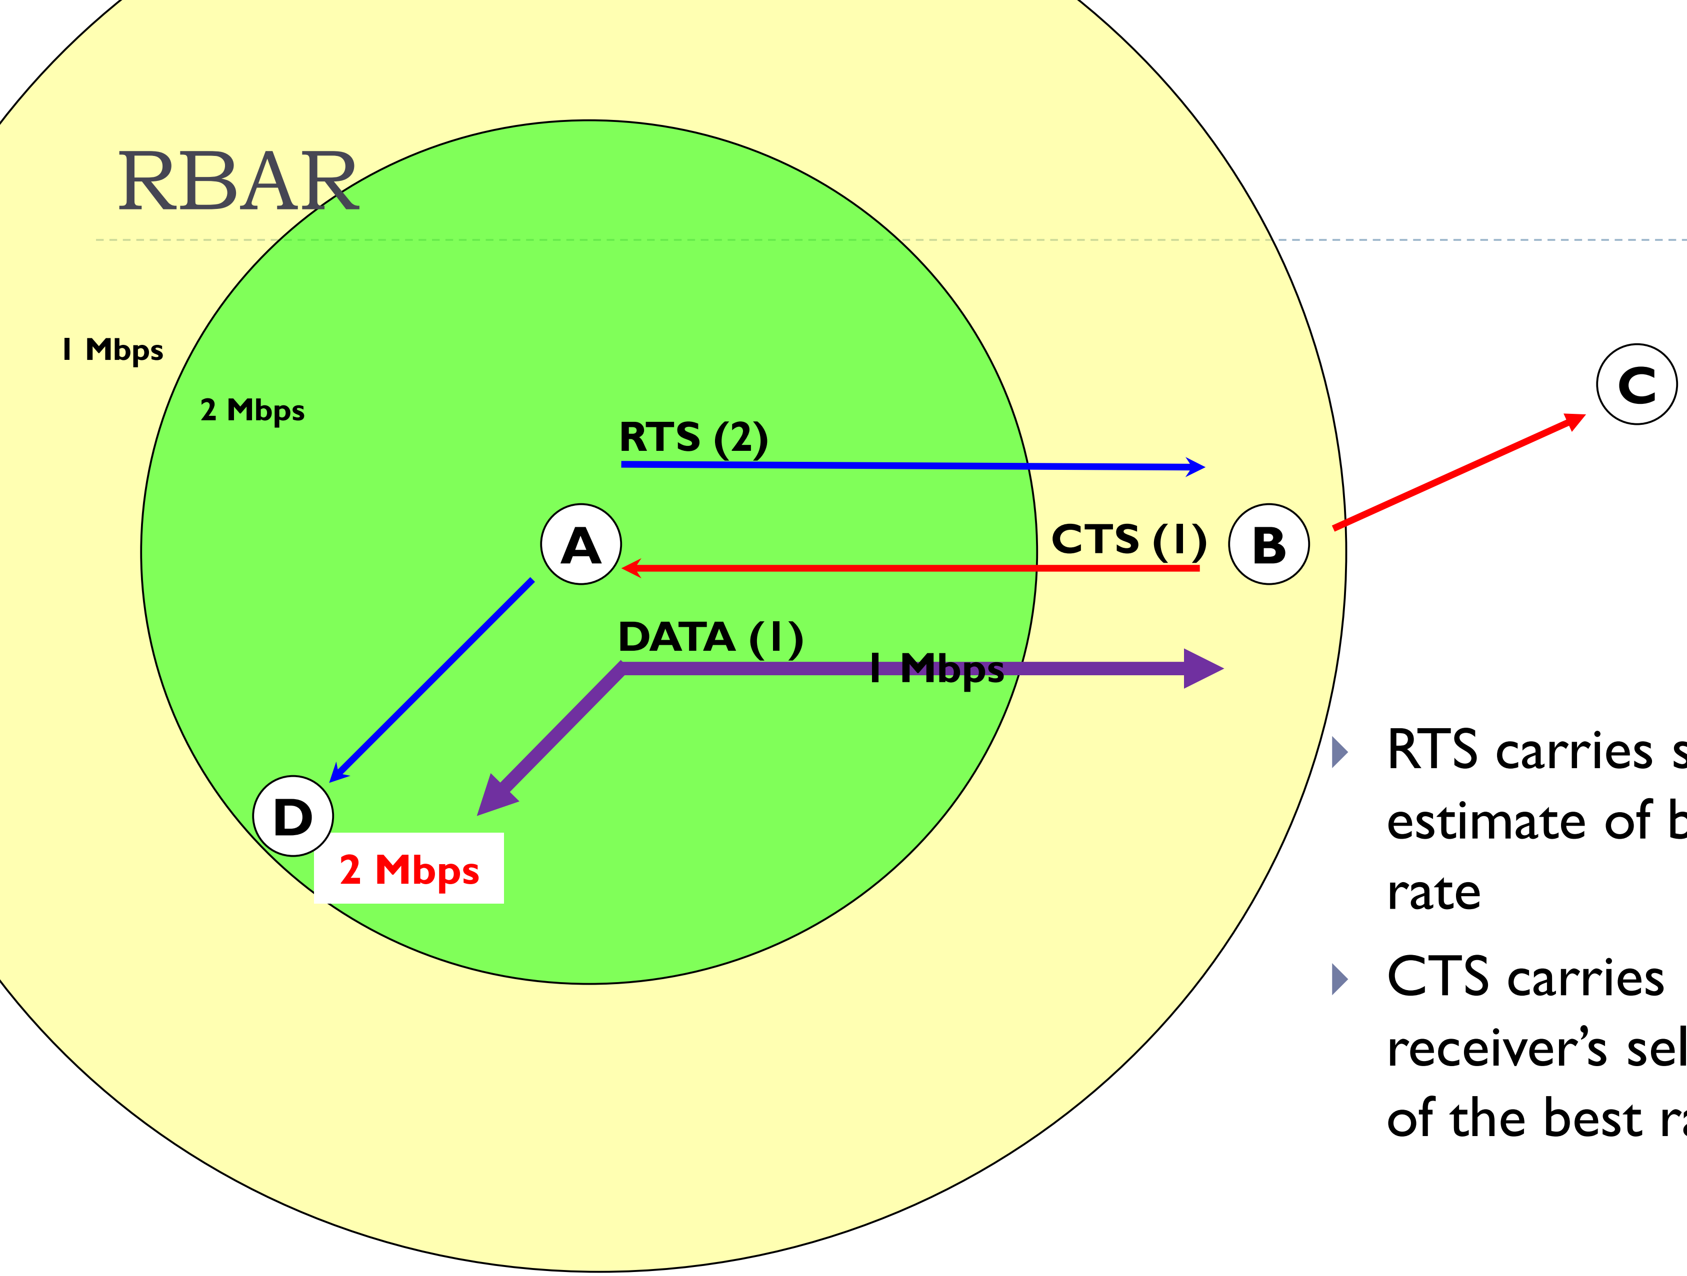

- RTS carries sender's estimate of best rate
- CTS carries receiver's selection of the best rate

RBAR

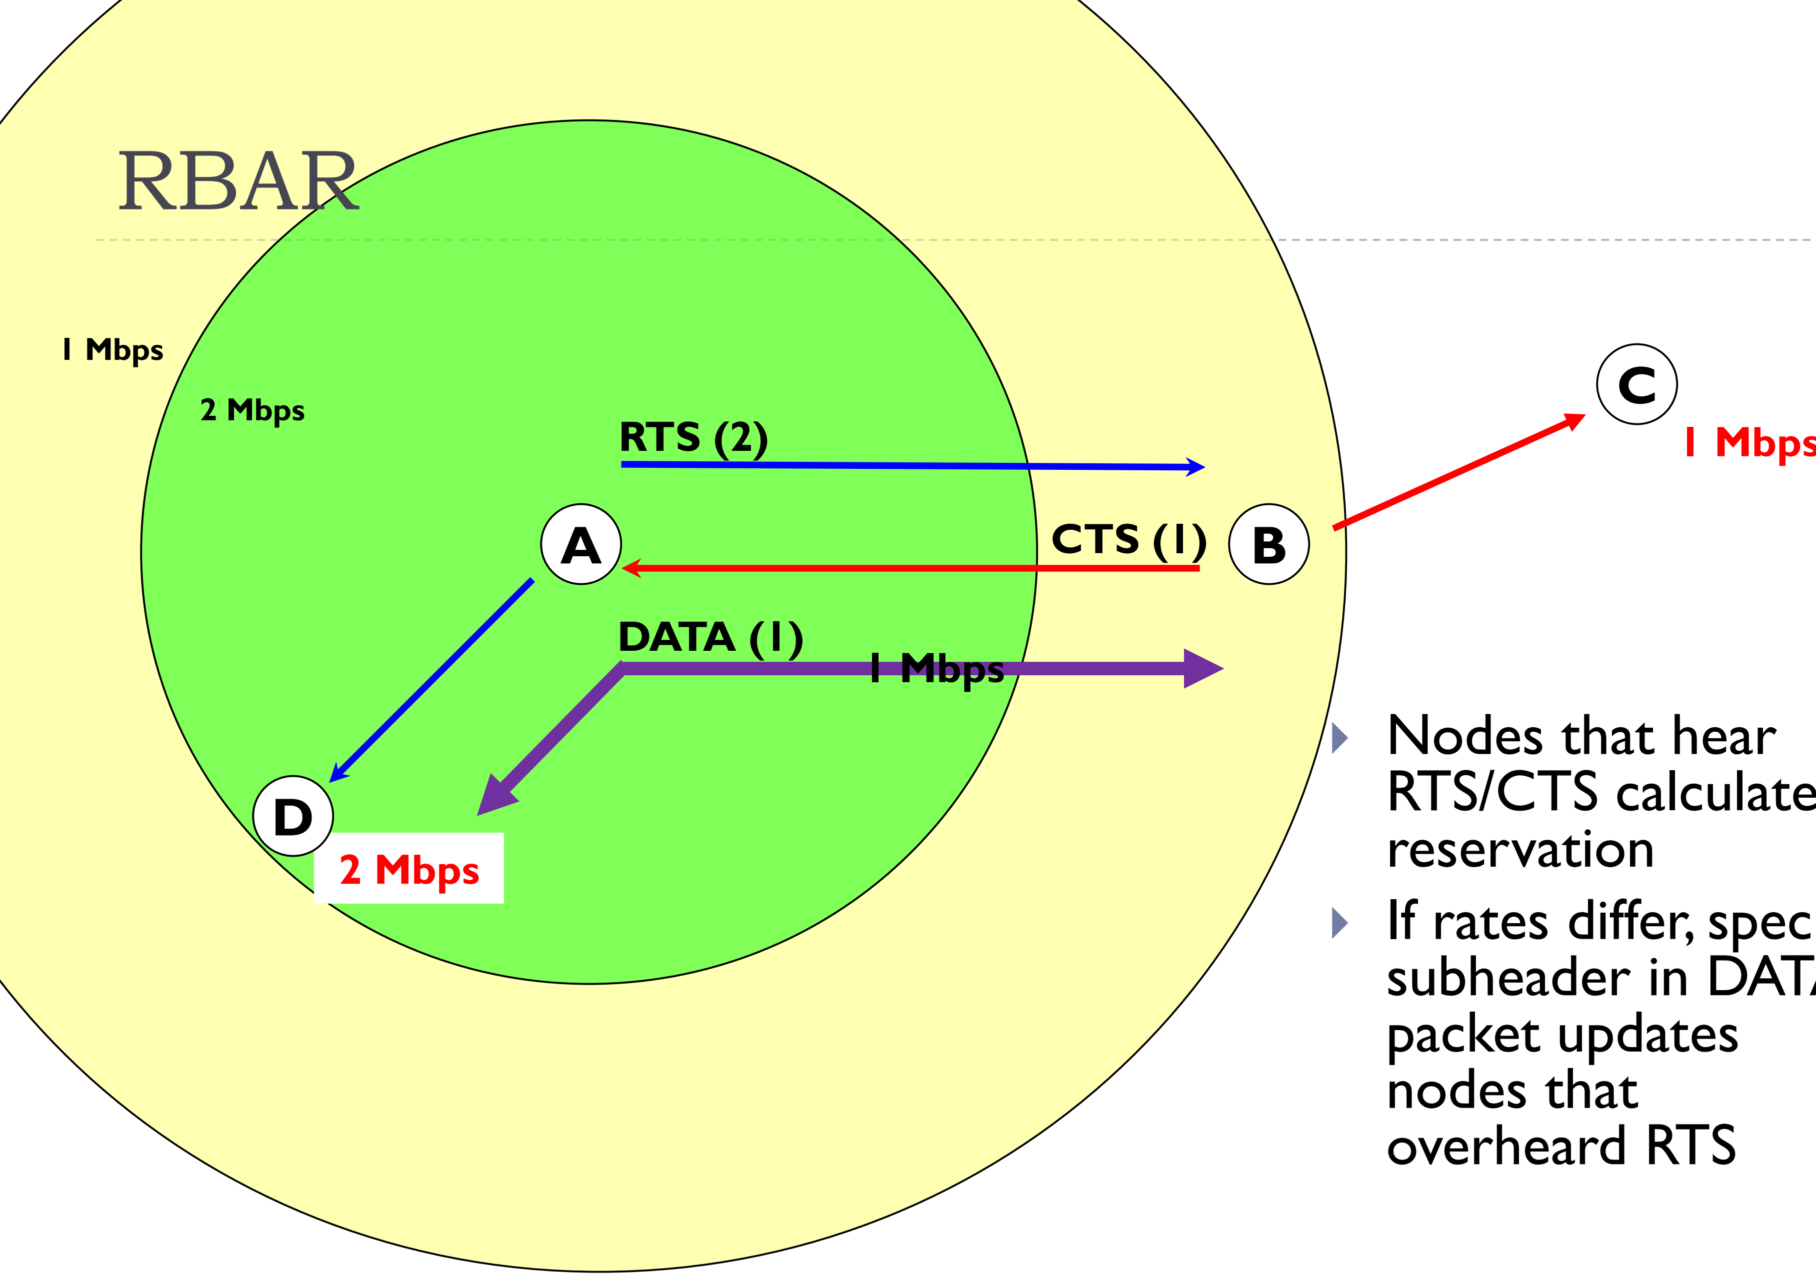

- Nodes that hear RTS/CTS calculate reservation
- If rates differ, special subheader in DATA packet updates nodes that overheard RTS

Performance of RBAR

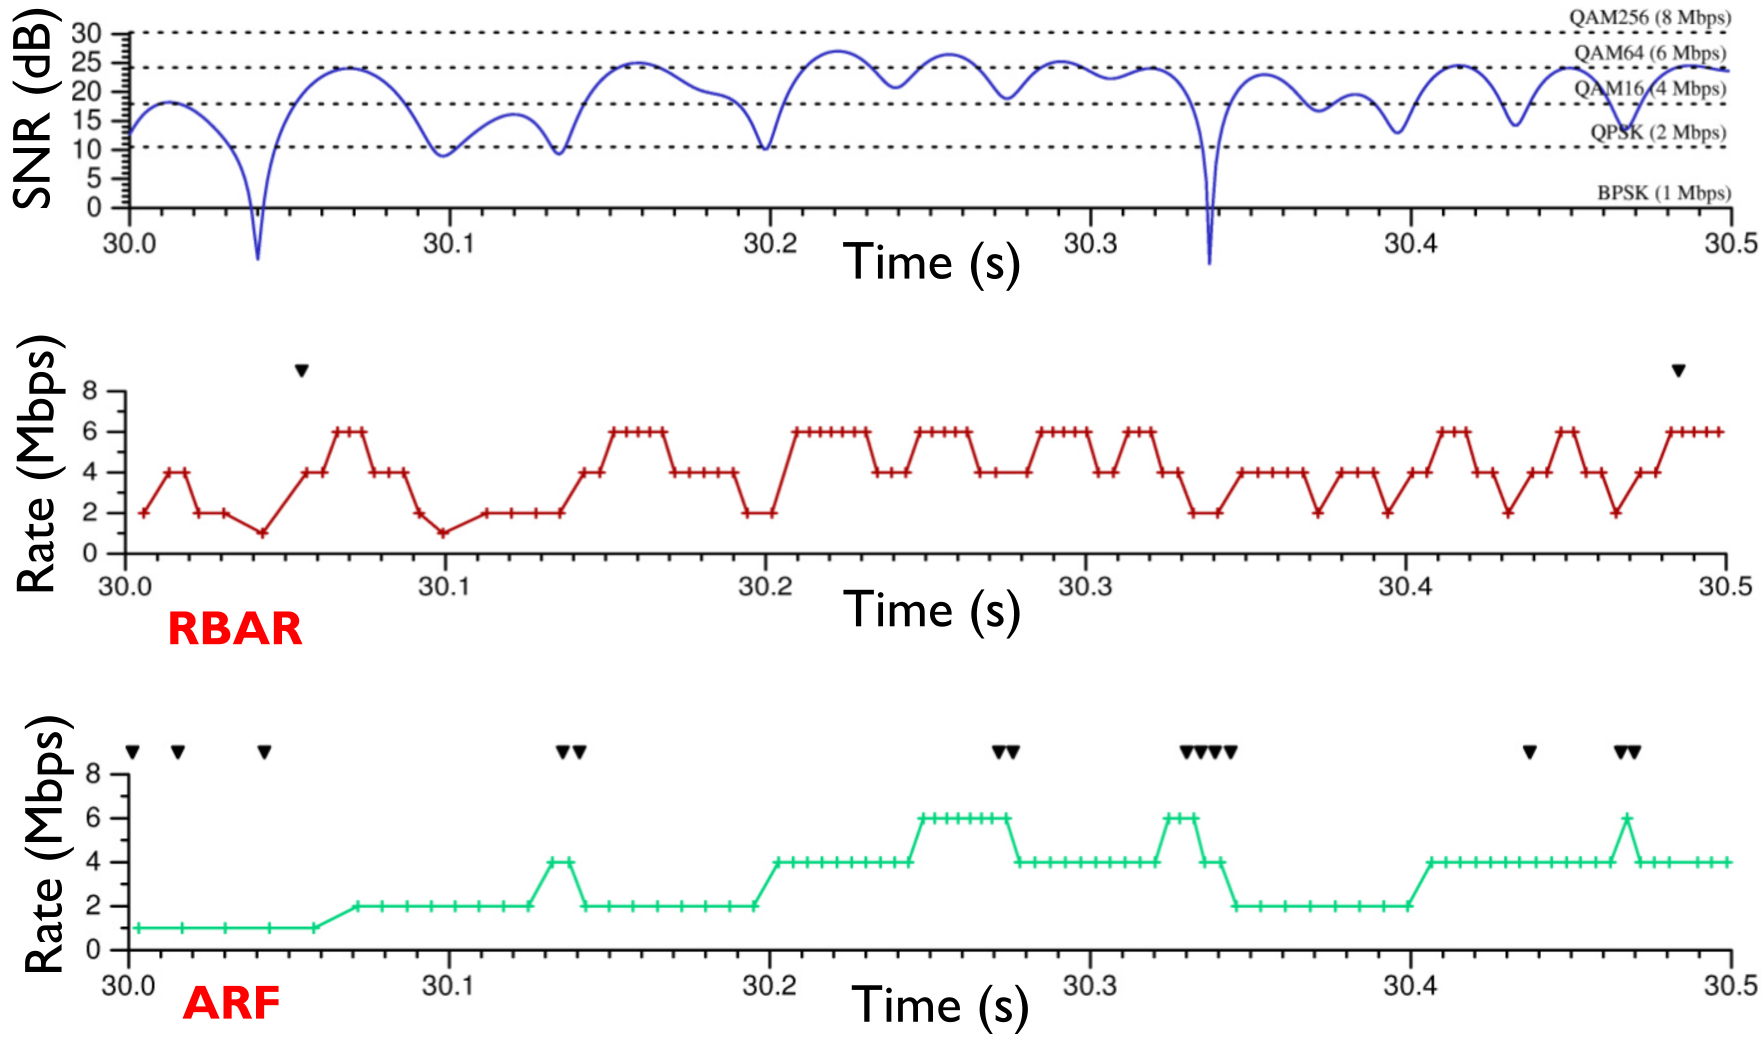

Implementation into 802.11

Frame Control	Duration	DA	SA	FCS	BSSID	Sequence Control	Body	FCS

⊢ Reservation Subheader (RSH) ⊣

- Encode data rate and packet length in duration field of frames
  - Rate can be changed by receiver
  - Length can be used to select rate
  - Reservations are calculated using encoded rate and length
- New DATA frame type with Reservation Subheader (RSH)
  - Reservation fields protected by additional frame check sequence
  - RSH is sent at same rate as RTS/CTS
- New frame is only needed when receiver suggests rate change

Evaluation

- Environment
  - Rayleigh fading
- Scenarios
  - Single-hop
- Protocols
  - RBAR and ARF

- RBAR
  - Channel quality prediction
    - SNR sample of RTS
  - Rate selection:
    - Threshold-based
  - Sender estimated rate:
    - Static (1 Mbps)

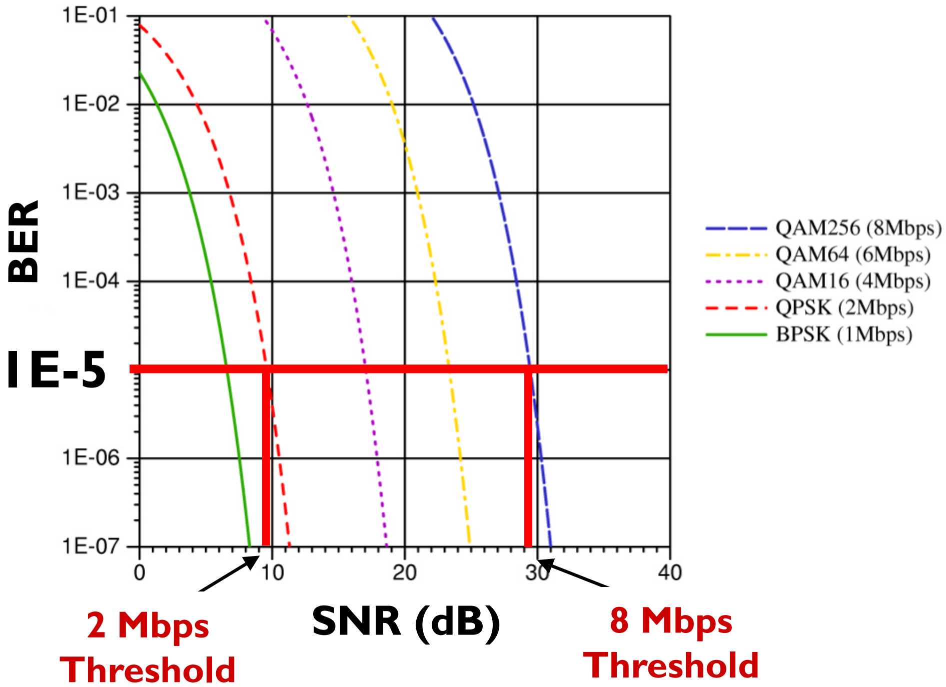

Single-Hop Scenario

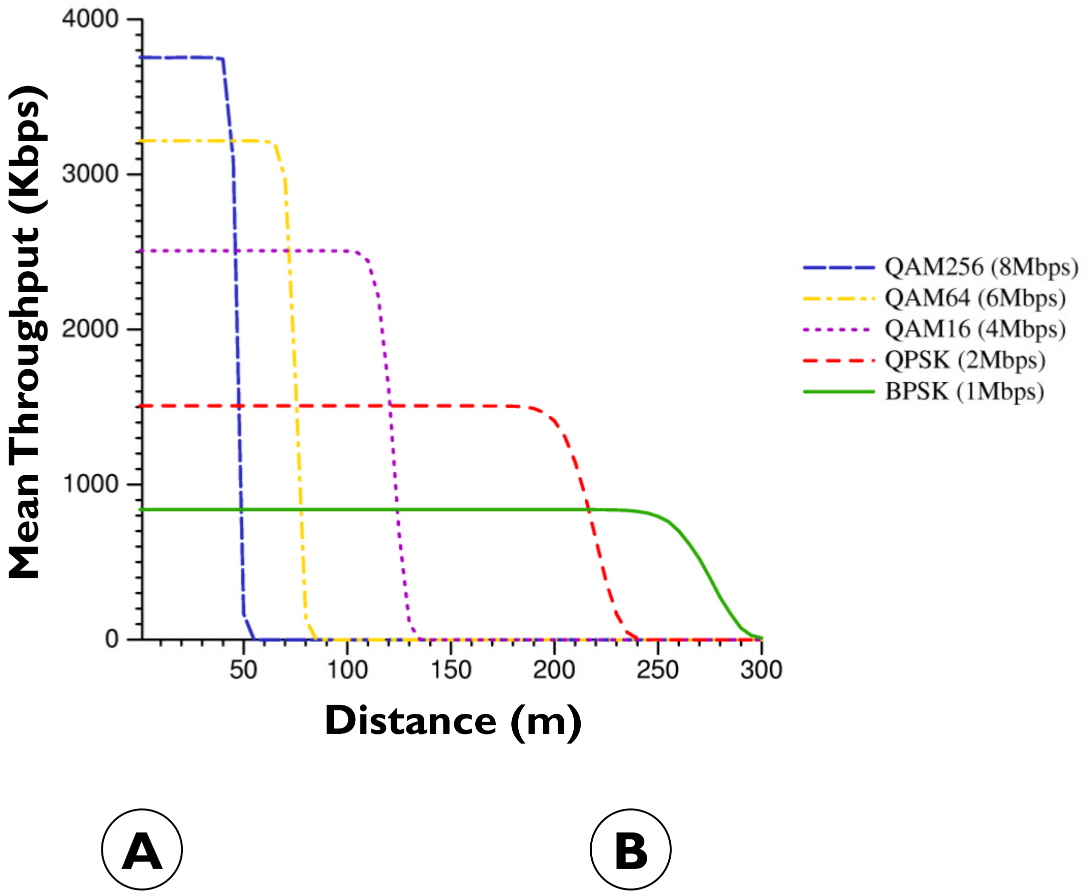

No Mobility - UDP Performance

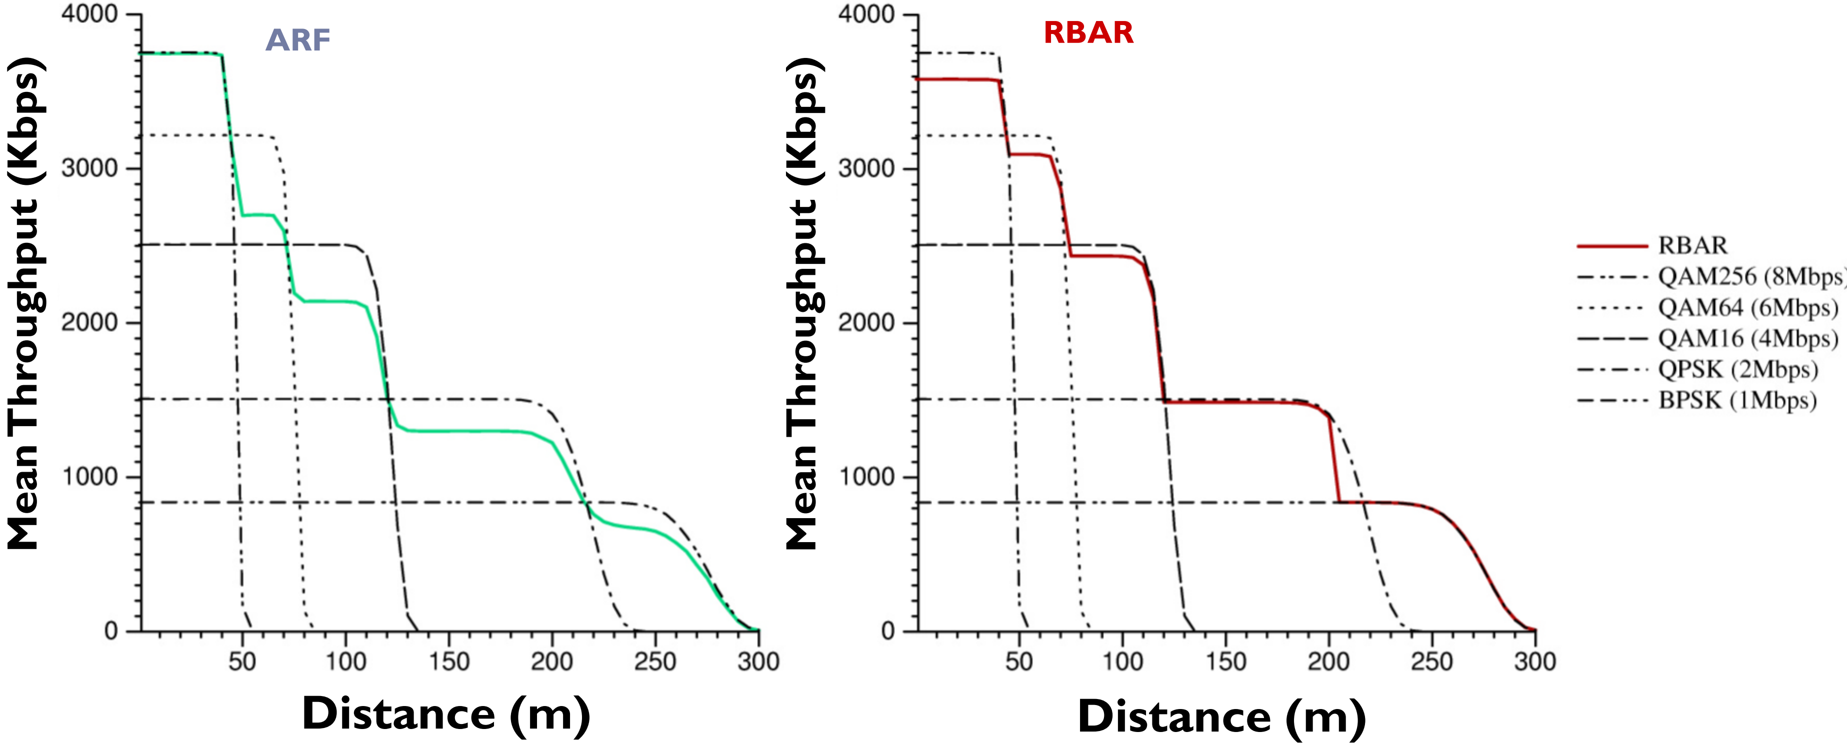

- RSH overhead seen at high data rates
  - Can be reduced using some initial rate estimation algorithm
- Limitations of simple threshold-based rate selection seen
- Generally, still better than ARF

No Mobility - UDP Performance

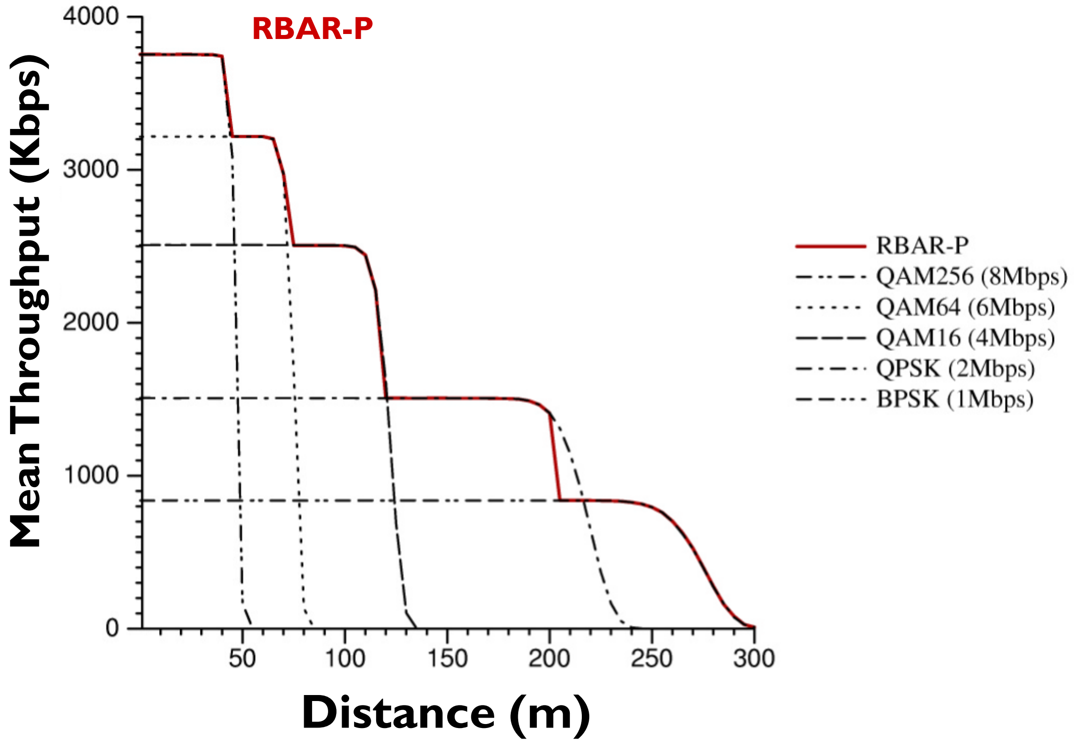

- RBAR-P – RBAR using a simple initial rate estimation algorithm
  - Previous rate used as estimated rate in RTS
- Better high-rate performance
- Other initial rate estimation and rate selection algorithms are a topic of future work

RBAR Summary

- Modulation schemes have different error characteristics
- Significant performance improvement may be achieved by MAC-level adaptive modulation
- Receiver-based schemes may perform best
  - Proposed Receiver-Based Auto-Rate (RBAR) protocol
  - Implementation into 802.11
- Future thoughts …
  - RBAR without use of RTS/CTS
  - RBAR based on the size of packets
  - Routing protocols for networks with variable rate links

Can we do better?

- Consider the situation below
  - ARF?
  - RBAR?

Motivation

- What if A and B are both at 56Mbps, and C is often at 2Mbps?
- Slowest node gets the most absolute time on channel?

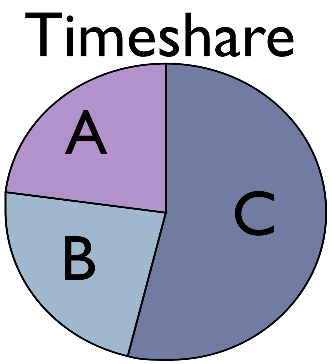

Throughput Fairness vs Temporal Fairness

MAC Layer Fairness Models

- Per Packet Fairness
  - If two adjacent senders continuously are attempting to send packets, they should each send the same number of packets
- Temporal Fairness
  - If two adjacent senders are continuously attempting to send packets, they should each be able to send for the same amount of medium time.
- In single rate networks these are the SAME!

Temporal Fairness Example

	802.11 Packet Fairness	OAR Temporal Fairness
11 Mbps Link	0.896	3.533
1 Mbps Link	0.713	0.450
Total Throughput	1.609	3.983

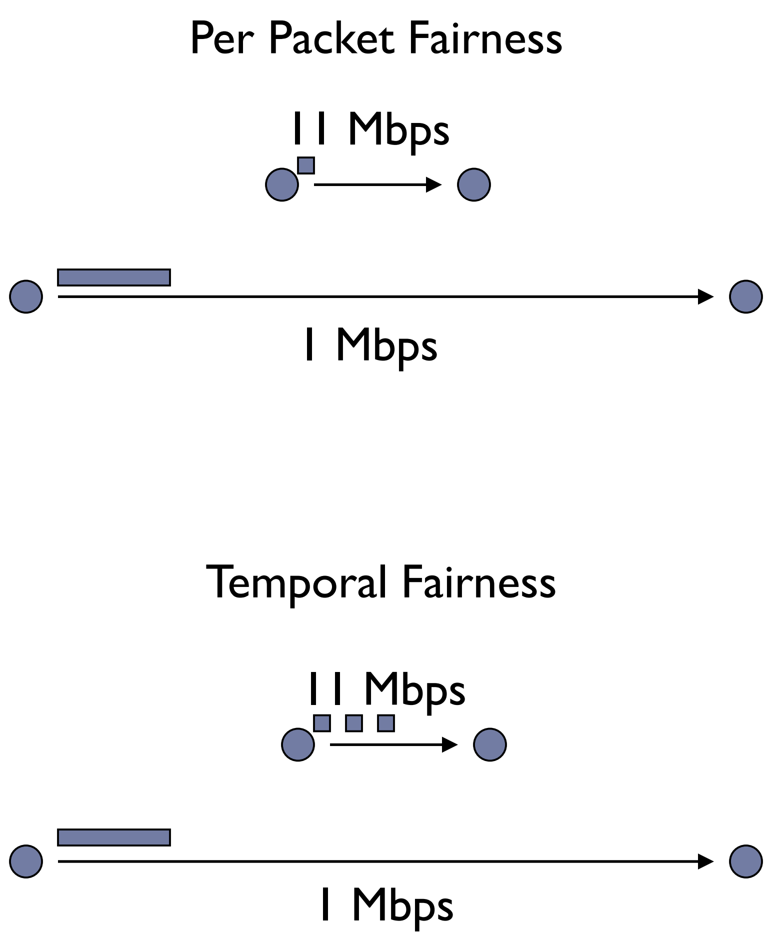

Opportunistic Scheduling

- Goal
  - Exploit short-time-scale channel quality variations to increase throughput
- Issue
  - Maintaining temporal fairness (time share) of each node
- Challenge
  - Channel info available only upon transmission

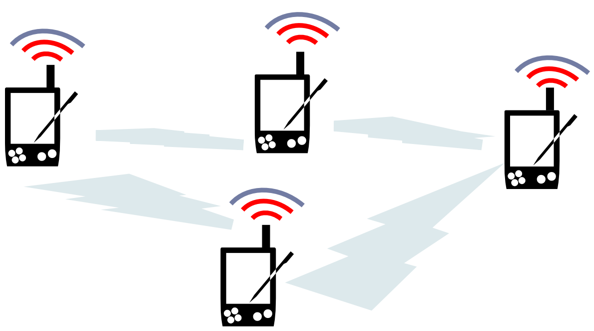

Opportunistic Auto-Rate (OAR)

- In many networks, there is intrinsic diversity
  - Exploiting this diversity can offer benefits
  - Transmit more when channel quality is high
    - else, free the channel quickly

- RBAR does not exploit this diversity
  - It optimizes per-link throughput

OAR Idea

- Basic Idea
  - Bad channel: transmit minimum number of packets
  - Good channel: transmit as much as possible

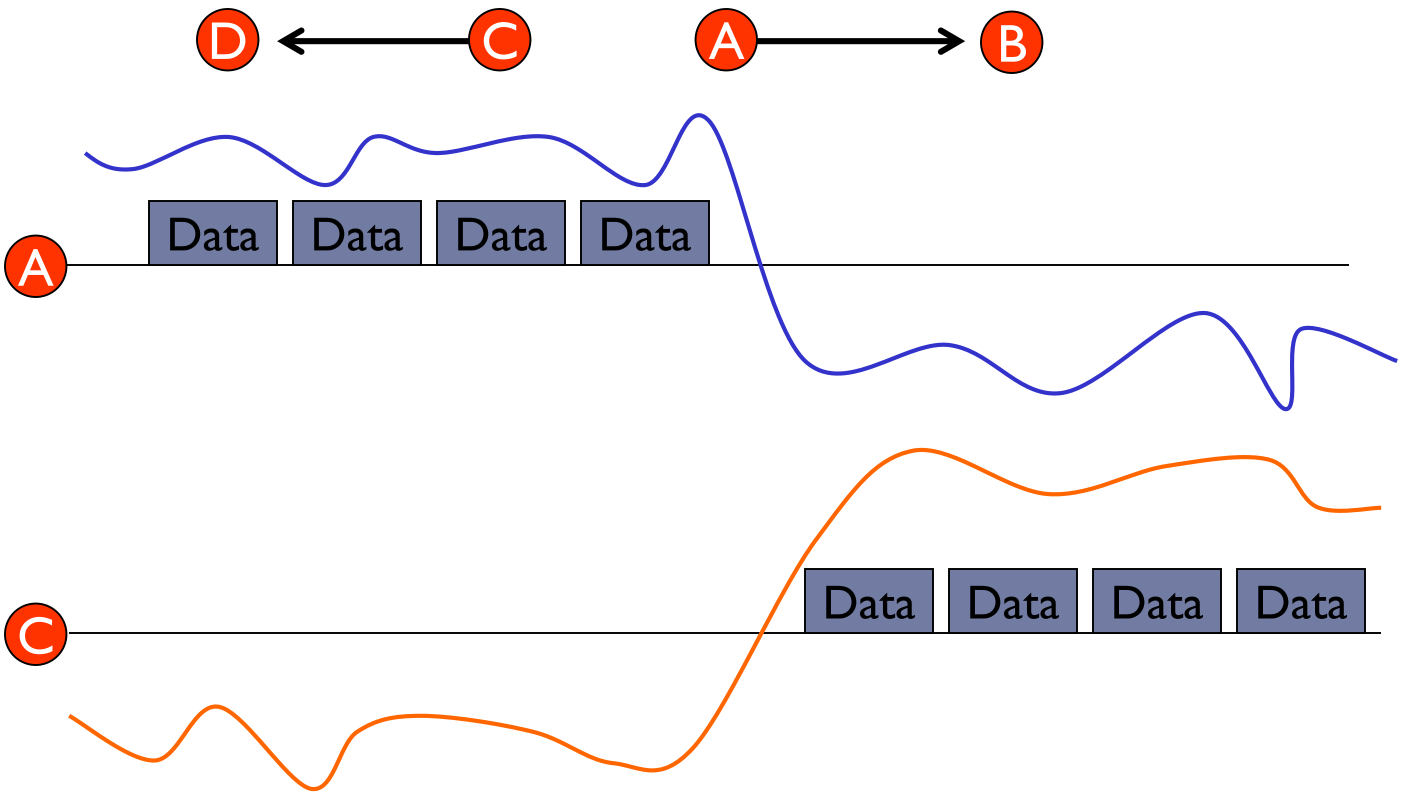

Why is OAR better?

- 802.11 alternates between transmitters A and C
  - Why is that bad

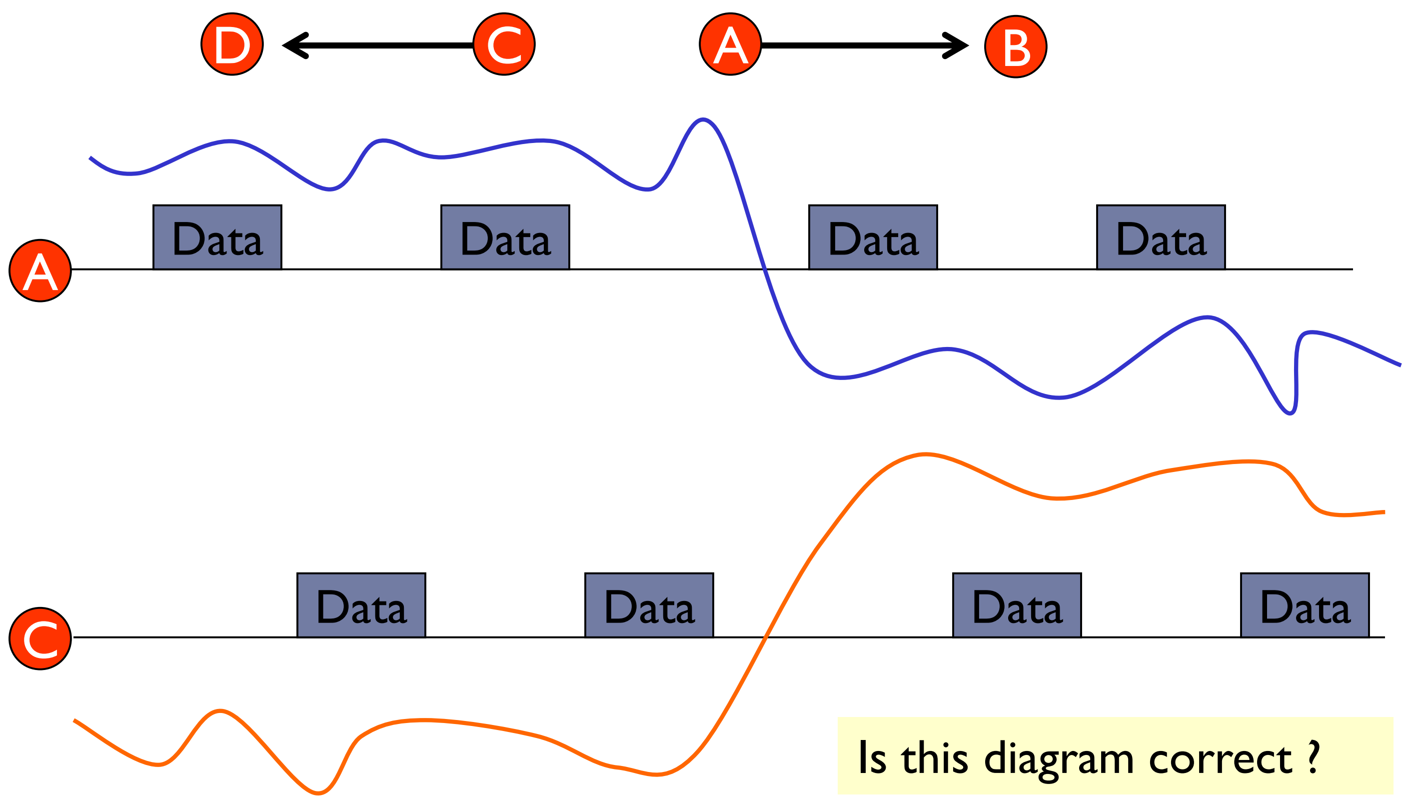

Why is OAR better ?

- Bad channel reduces SINR → increases transmit time
  - Fewer packets can be delivered

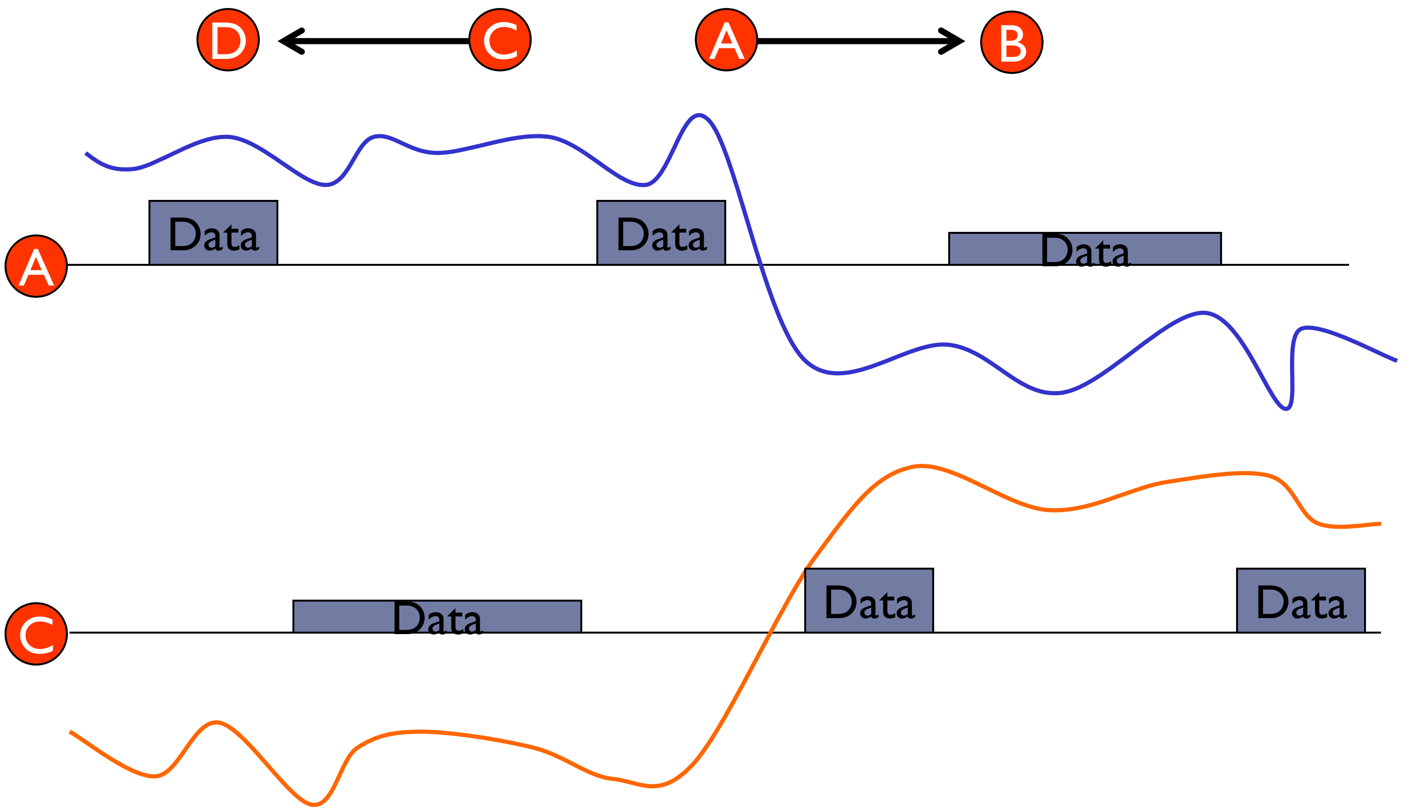

OAR Protocol Steps

- Transmitter estimates current channel
  - Can use estimation algorithms
  - Can use RBAR, etc.
- If channel better than base rate (2 Mbps)
  - Transmit proportionally more packets
    - e.g., if channel can support 11 Mbps, transmit (11/2 ~ 5) pkts
- OAR upholds temporal fairness
  - Each node gets same duration to transmit
  - Sacrifices throughput fairness → the network gains!!

OAR Protocol

Protocol	Channel Condition					
	BAD		MEDIUM		GOOD	
	Pkts	Rate	Pkts	Rate	Pkts	Rate
802.11	1	2	1	2	1	2
802.11b	1	2	1	5.5	1	11
OAR	1	2	3	5.5	5	11

- Rates in IEEE 802.11b: 2, 5.5, and 11 Mbps

Evaluation

- Simulation experiments
  - Fully connected network: all nodes in radio range of each other
    - Number of Nodes, channel condition, mobility, node location

Fully Connected Setup

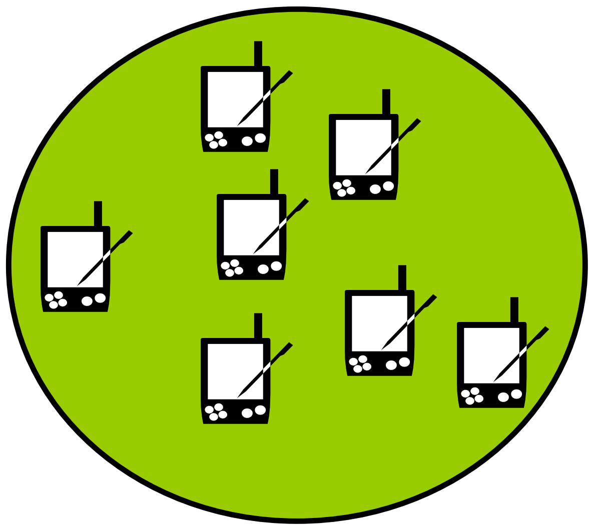

- Every node can communicate with everyone
- Each node's traffic is at a constant rate and continuously backlogged
- Channel quality is varied dynamically

Fully Connected Throughput Results

- OAR vs. RBAR
  - 42% to 56% gain
- Gain increases with the number of flows
- Note
  - Both RBAR and OAR are significantly better than standard 802.11
  - 230% and 398% respectively

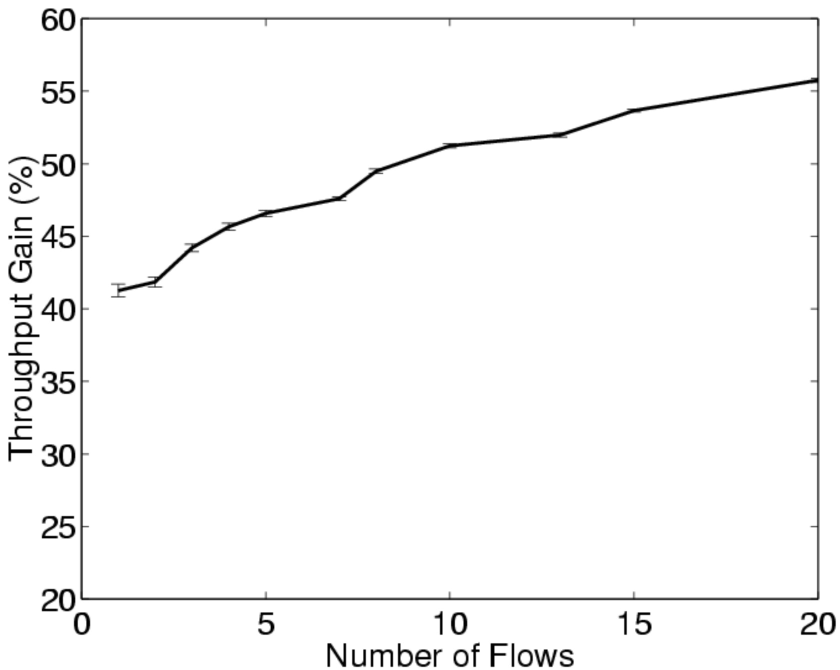

OAR thoughts

- OAR does not offer benefits when
  - Neighboring nodes do not experience diverse channel conditions
  - Coherence time is shorter than N packets

Summary

- Rate control can be useful
  - When adapted to channel fluctuations (RBAR)
  - When opportunistically selecting transmitters (OAR)
- Benefits maximal when
  - Channel conditions vary widely in time and space
- Correlation in fluctuation can offset benefits
  - OAR may show negligible gains

What lies ahead ?

- Dual of rate-control is power control
  - One might be better than the other
  - Decision often depends on the scenario → open problem