2. Other DCF features, limits and extensions

Fragmentation

- ➔ **Splits message (MSDU) into several frames (MPDU)**
  - ⇨ Same fragment size
    - → except the last one
- ➔ **Fragmentation burst**
  - ⇨ Fragments separated by SIFS
    - → Channel cannot be captured by someone else
  - ⇨ Each fragment individually ACKed
- ➔ **Each fragment reserves channel for next one**
  - ⇨ NAV updated fragment by fragment
- ➔ **Missing ACK for fragment x**
  - ⇨ Release channel (automatic)
  - ⇨ Backoff
  - ⇨ Restart from transmission of fragment x

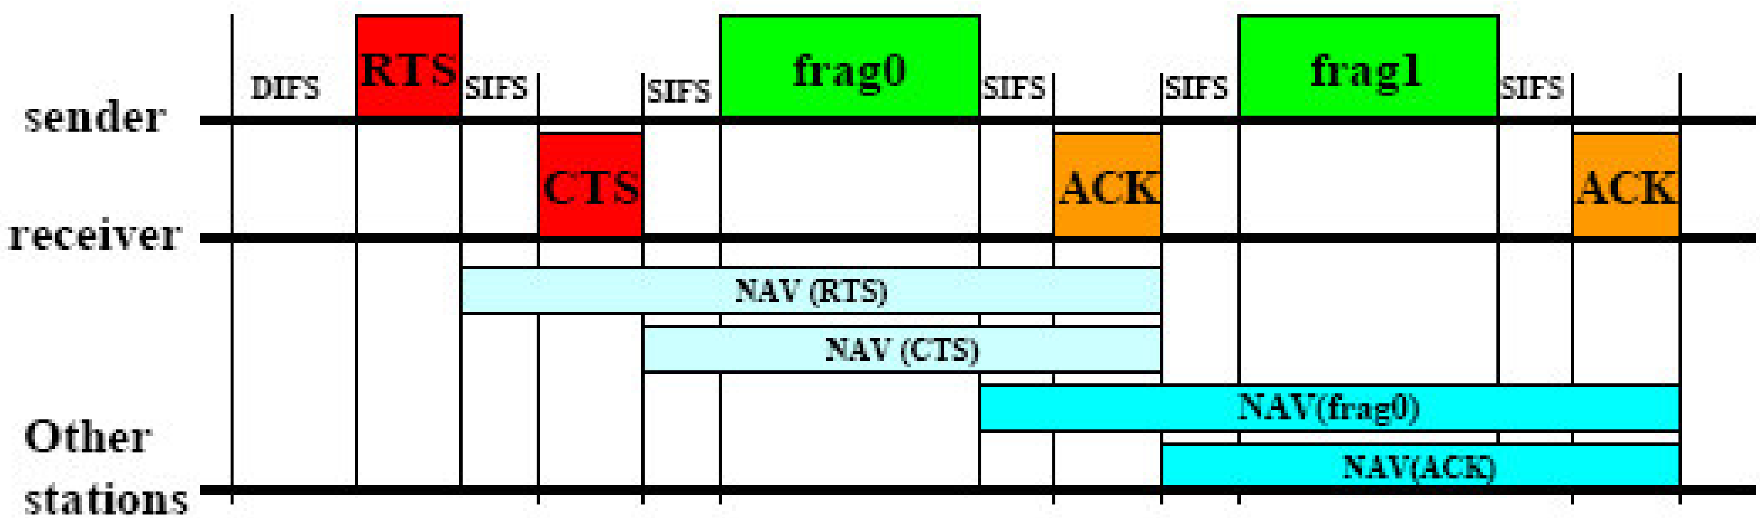

Why Fragmentation?

- ➔ **High Bit Error Rate (BER)**
  - ⇨ Increases with distance
  - ⇨ The longer the frame, the lower the successful TX probability
  - ⇨ High BER = high rts overhead & increased rtx delay
    - → **Backoff window increases: cannot distinguish collisions from tx error!**

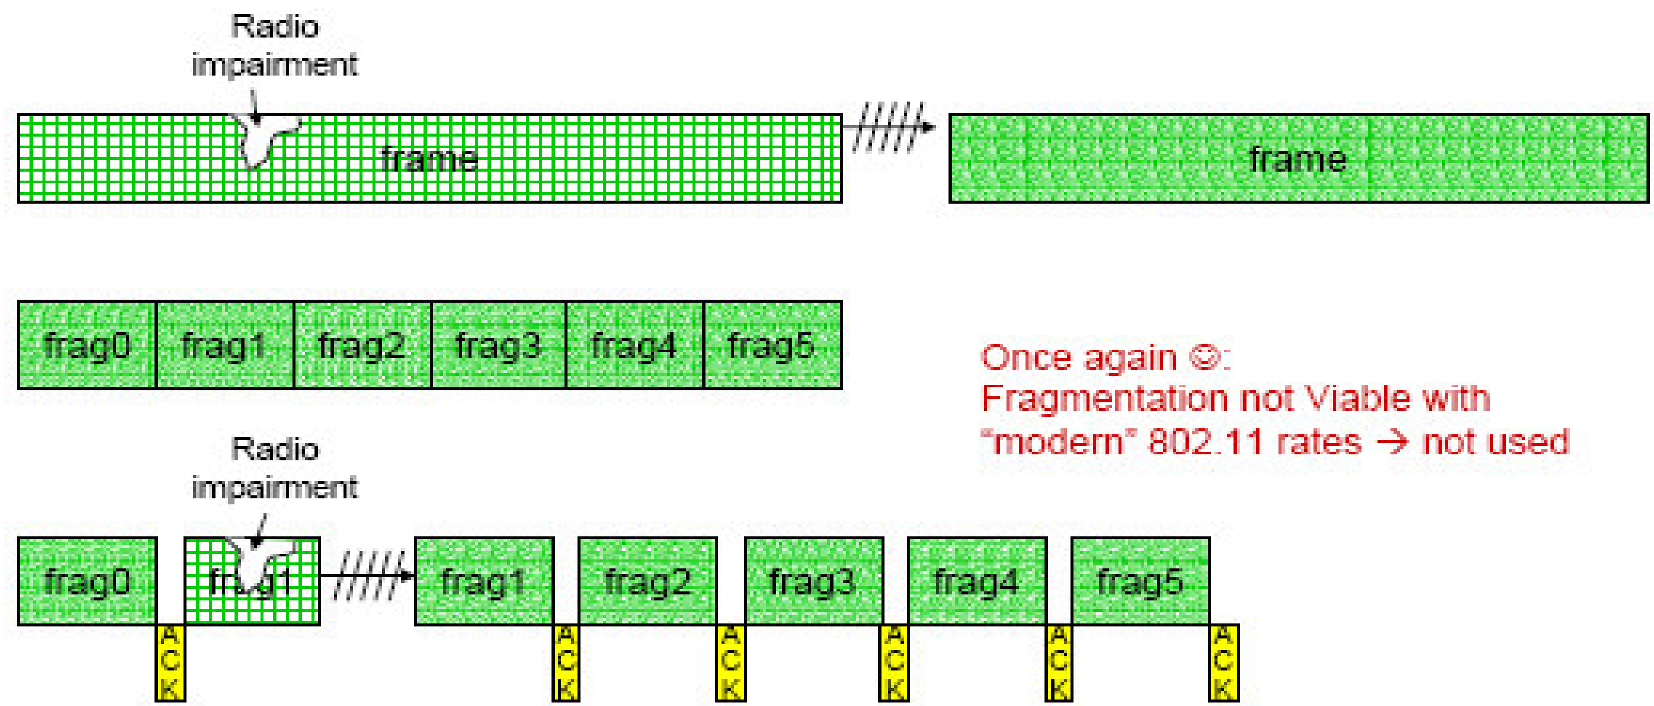

Once again ☺:
Fragmentation not Viable with "modern" 802.11 rates → not used

Fragment and sequence numbers

DATA FRAME (28 bytes excluded address 4)

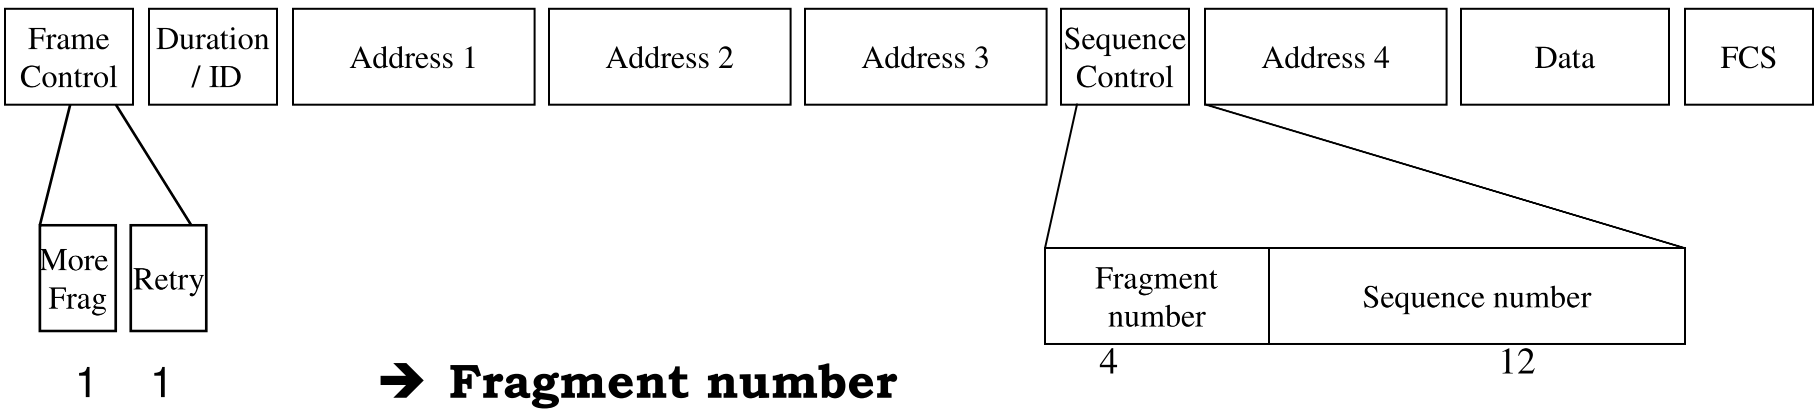

- ➔ **Fragment number**
  - ⇨ Increasing integer value 0-15 (max 16 fragments since 4 bits available)
  - ⇨ Essential for reassembly
- ➔ **More fragment bit (frame control field) set to:**
  - ⇨ 1 for intermediate fragments
  - ⇨ 0 for last fragment
- ➔ **Sequence Number**
  - ⇨ Used to filter out duplicates
    - → Unlike Ethernet, duplicates are quite frequent!
    - → Retransmissions are a main feature of the MAC
- ➔ **Retry bit: helps to distinguish retransmissions**
  - ⇨ Set to 0 at transmission of a new frame

Multi-rate operation

- ➔ **Rate selection: proprietary mechanism!**
  - ⇨ Result: different chipsets operate widely different
- ➔ **Two basic approaches**
  - ⇨ Adjust rate according to measured link quality (SNR estimate)
    - → How link quality is computed is again proprietary!
  - ⇨ Adjust rate according to frame loss
    - → How many retries? Step used for rate reduction?
    - → Problem: large amount of collisions (interpreted as frame loss) forces rate adaptation

Performance Anomaly

[M. Heusse, et al. "Performance Anomaly of 802.11b", INFOCOM 2003]

- **Question 1:**
  - Assume that throughput measured for single 11 mbps greedy stations is approx 6 mbps. What is pr-STA throughput when two 11 mbps greedy stations compete?
- **Answer 1:**
  - Approx 3 mbps (easy!
  -)
- **Question 2:**
  - Assume that throughput measured for a single 2 mbps greedy stations is approx1.7 mbps. What is per-STA throughput when two 2 mbps greedy stations compete?
- **Answer 2:**
  - Approx 0.85 mbps (easy!)

- **Question 3:**
  - What is the per-STA throughput when one 11 mbps greedy station compete with one 2 mbps greedy station?
- **Answer 3:**
  - …

An intuitive answer..

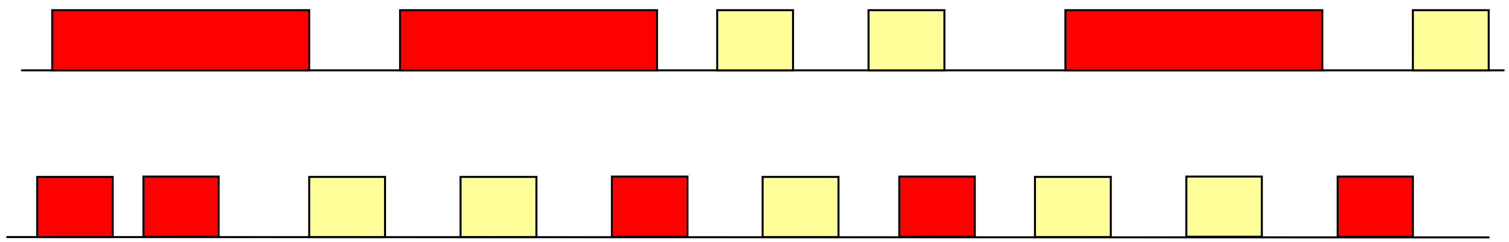

- ➔ **The probability that at each contention a given station gets the next channel access (i.e. extracts the lower backoff) is fixed for all the stations!**
- ➔ **In long terms, all the stations receive the same number of transmission grants**
  - ⇨ If payload size is fixed: the throughput of high rate and low rate stations is the same, regardless of the transmission rate
    - →***throughput fairness***
    - →low rate stations waste resources for high rate stations

Understanding Answers 1&2
(neglect collisions – indeed rare with only two stations)

In average, STA1 and STA2 alternate their transmissions on the channel!

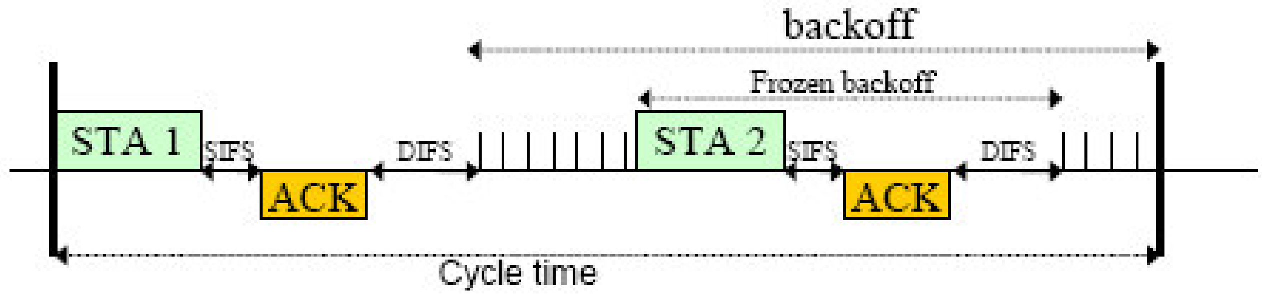

$$Thr[1] = Thr[2] = \frac{E[payload]}{E[cycle\ time]} = \frac{1500 \times 8}{T_{MPDU}[1] + SIFS + ACK + DIFS + T_{MPDU}[2] + SIFS + ACK + DIFS + E[backoff]}$$

- **Data Rate = 11 mbps; ACK rate = 1 mbps**
- **Payload = 1500 bytes**

$$T_{MPDU} = 192 + 8 \cdot (28 + 1500)/11 \approx 1303$$
$$T_{ACK} = 192 + 8 \cdot 14/1 = 304$$
$$SIFS = 10; \quad DIFS = 50$$
$$E[Backoff] = \frac{31}{2} \times 20 = 310$$
$$Thr = \frac{1500 \times 8}{2 \times (1303 + 10 + 304 + 50) + 310} = 3.3 Mbps$$

- **Data Rate = 2 mbps; ACK rate = 1 mbps**
- **Payload = 1500 bytes**

$$T_{MPDU} = 192 + 8 \cdot (28 + 1500)/2 \approx 6304$$
$$T_{ACK} = 192 + 8 \cdot 14/1 = 304$$
$$SIFS = 10; \quad DIFS = 50$$
$$E[Backoff] = \frac{31}{2} \times 20 = 310$$
$$Thr = \frac{1500 \times 8}{2 \times (6304 + 10 + 304 + 50) + 310} = 0.88 Mbps$$

Computing answer 3

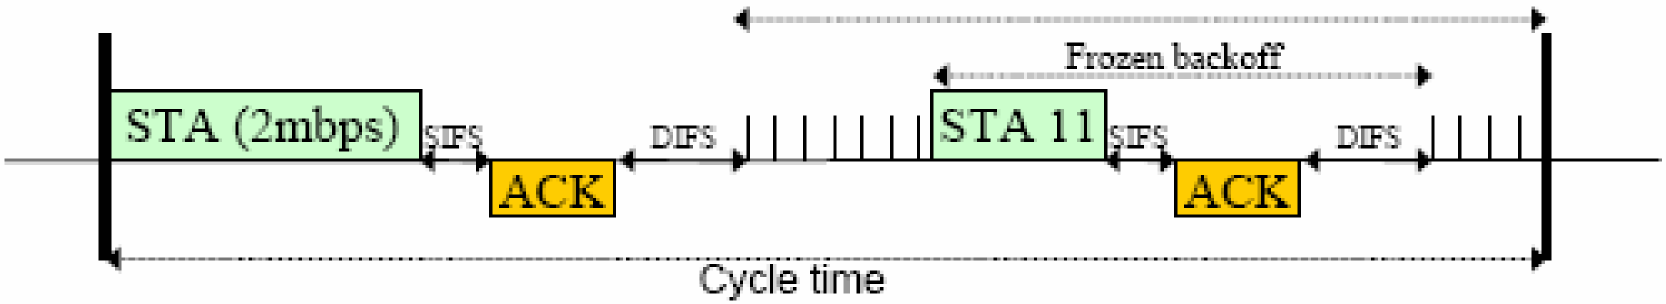

RESULT: SAME THROUGHPUT (in the long term)!!

$$Thr[1] = Thr[2] = \frac{E[payload]}{E[cycle\ time]} =$$

$$= \frac{1500 \times 8}{T_{MPDU}[1] + SIFS + ACK + DIFS + T_{MPDU}[2] + SIFS + ACK + DIFS + E[backoff]} =$$

$$= \frac{1500 \times 8}{6304 + 1303 + 2(10 + 304 + 50) + 310} = 1.39\ Mbps!!!!!!$$

DRAMATIC CONSEQUENCE: throughput is limited by STA with slowest rate (lower that the maximum throughput achievable by the slow station)!!

Performance anomaly into action

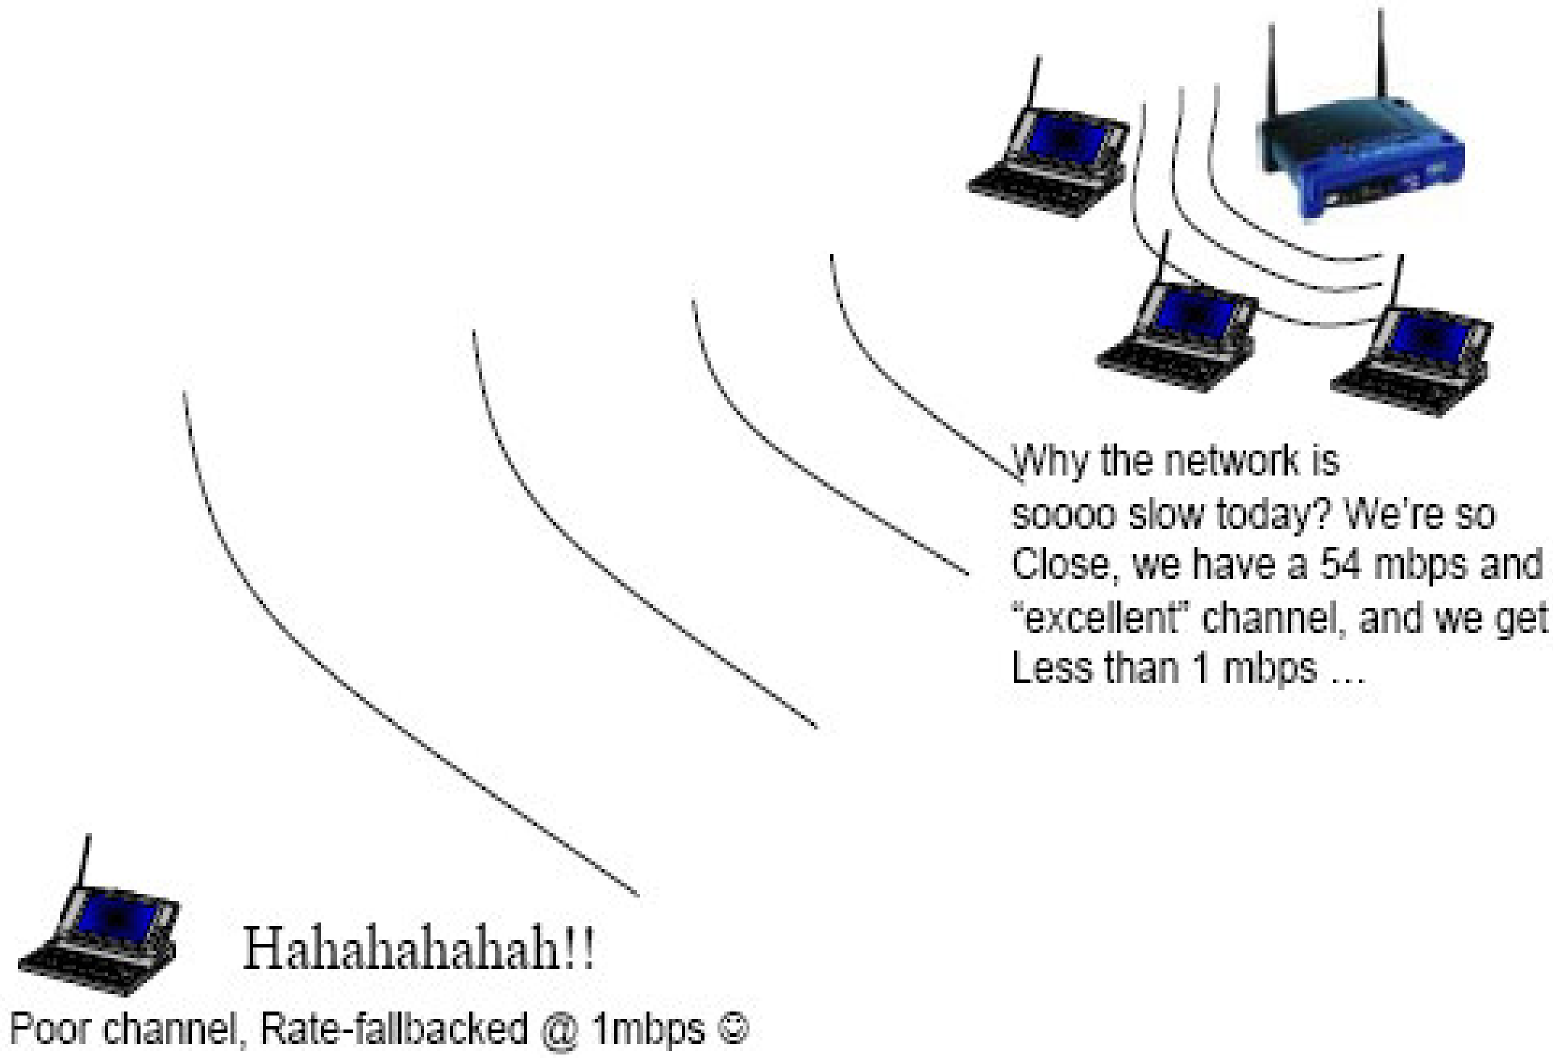

Spatial reuse

Transmission/Interference/CS Range

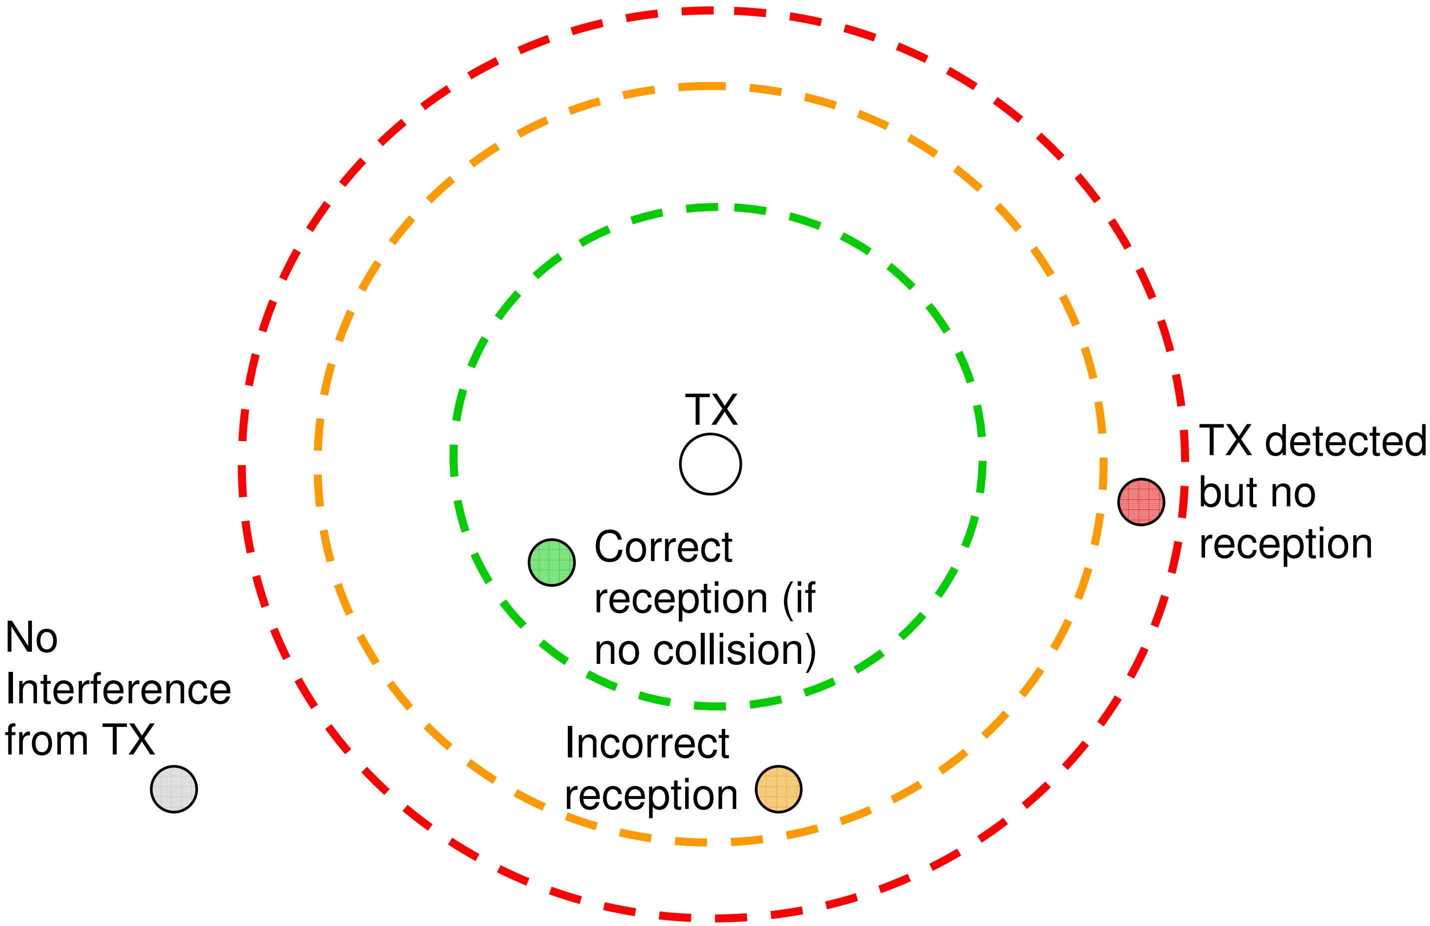

Exposed Nodes

- ➔ Any node within carrier sense range of transmitter and out of interference range of receiver
- ➔ Prevents simultaneous transmissions
- ➔ Reduction in Spatial Reuse
  - ⇨ **C in carrier sense range of A && out of interference range of B**

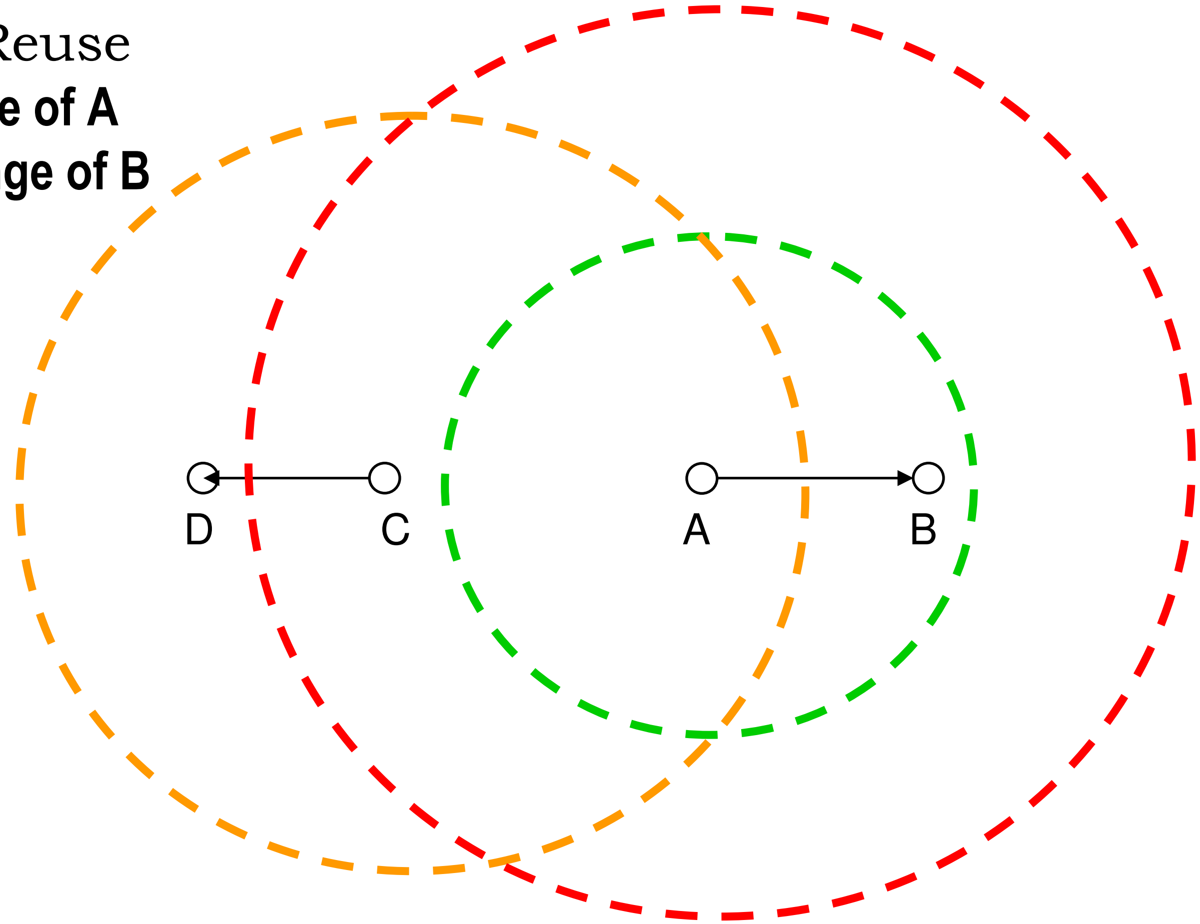

Is exposed node a problem?

- ➔ Not really!
- ➔ Remember that DCF handshake is asynchronous…
  - ⇨ **If C tx to D && a tx to B,**
    No interference @ D & B
    BUT:
    C still TX to D && B replies to A
    with an ACK ->
    Interference on A!!!!

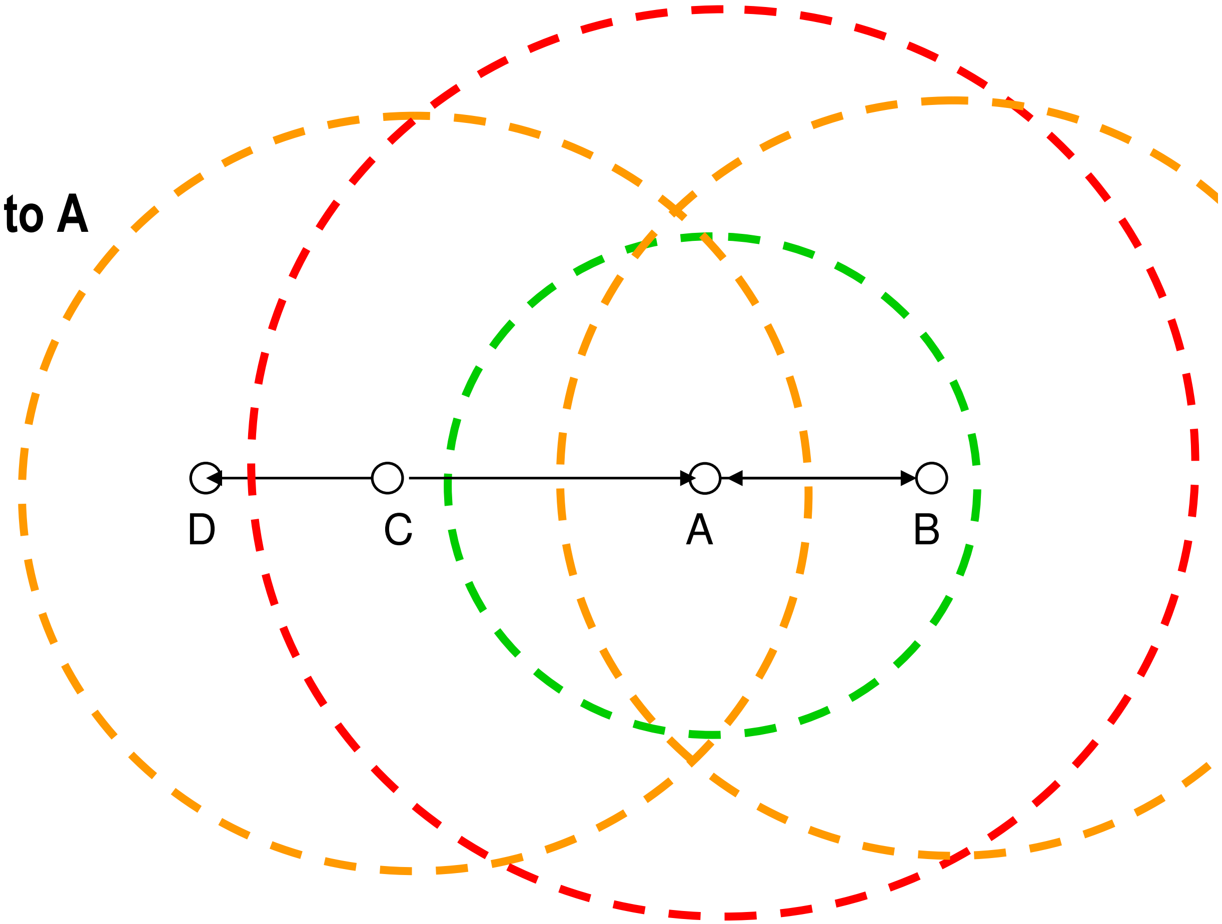

Node chains

- **In practical scenarios, packets can be often delivered from source to destination through multiple radio hops**
- **dramatic performance impairment in node chains**
  - Nodes can forward only a single packet at a time, blocking neighbor transmissions
  - Hidden nodes

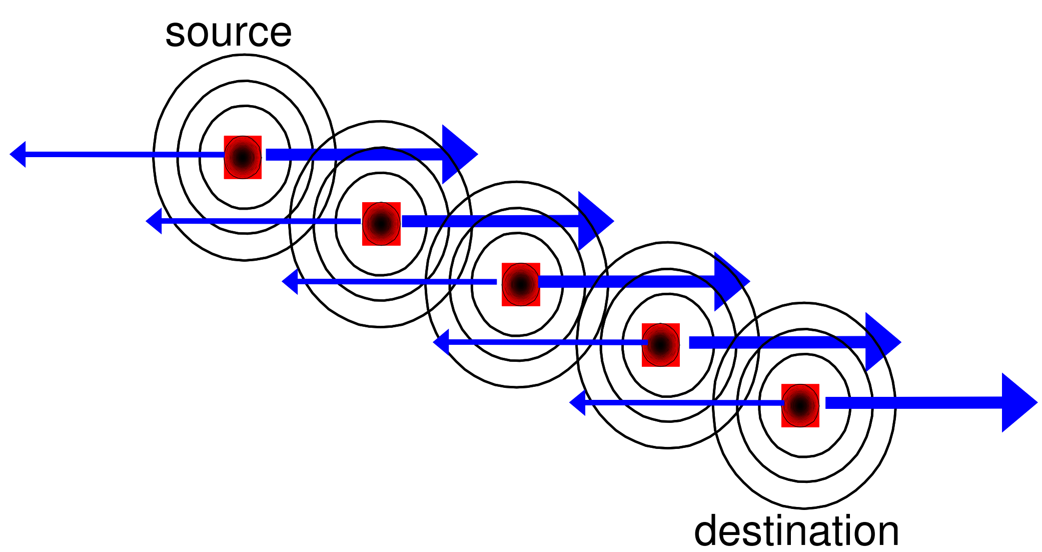

Chain capacity
[J. Li, et al "Capacity of Ad Hoc Wireless Networks"]

- ➔ **Assume that Transmission, Interference and CS ranges coincide**
- ➔ **Simultaneous transmissions along the chain:**
  - ⇨ If node distance = CS+1 -> collision! (e.g. back collision at node 2!)
  - ⇨ If node distance > CS+1 -> spatial reuse. (e.g. node 4 and node 7 receive correctly!)

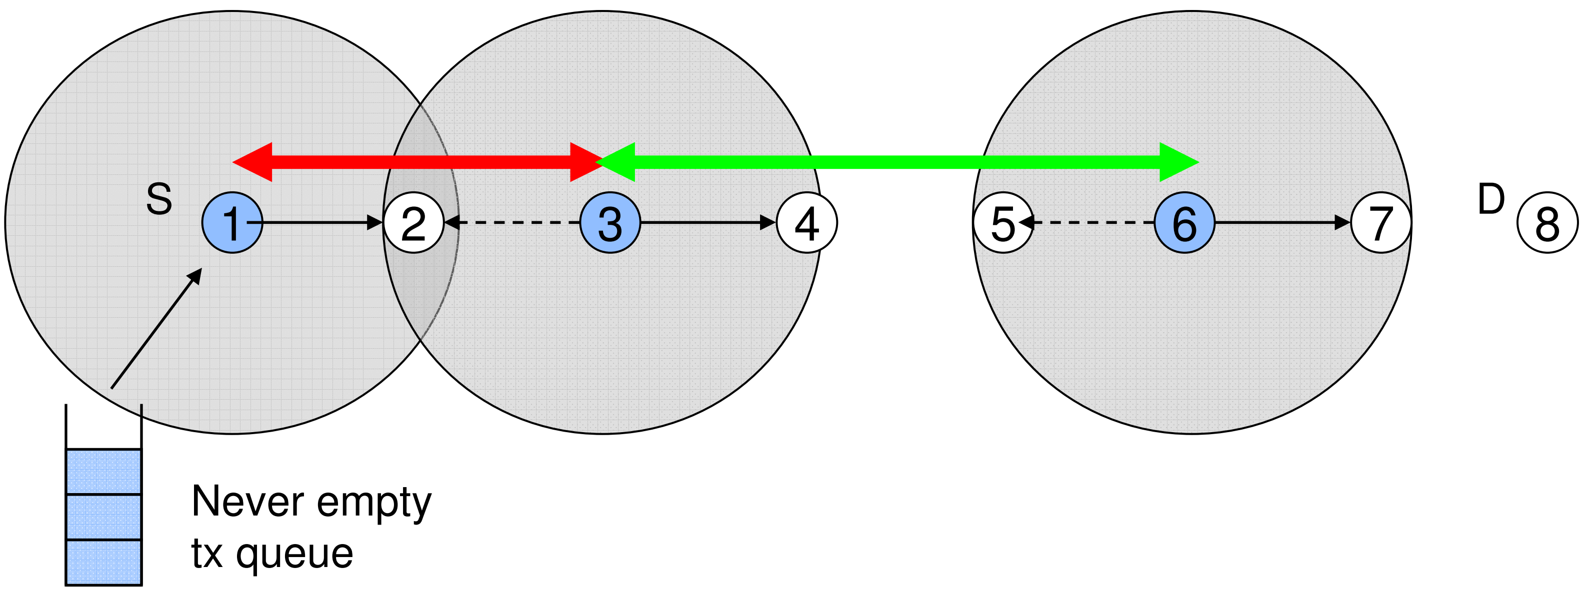

Maximum chain capacity

- Question: if *r* is the throughput when node 1 transmits alone towards node 2, what is the maximum packet delivery rate between 1 and 8, assuming ideal packet scheduling?
- Answer: if tx order is 1-2-3-4-5-6-7-8, we have 7 tx before a packet delivery -> max thr = r/7
- Answer: we can exploit simultaneous tx! After a transient tx order cyclically is (1,4,7)-(2,5)-(3,6): we have 3 tx before a packet delivery -> max thr = r/3

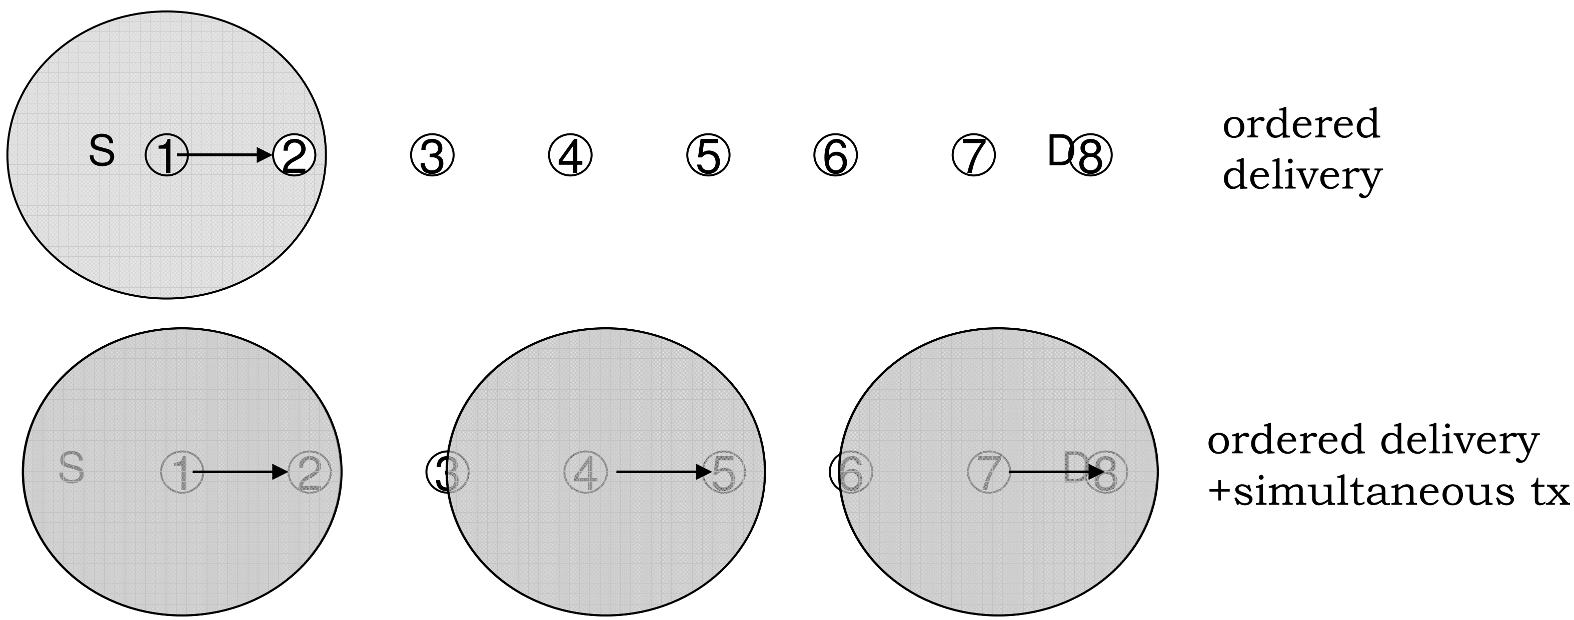

Maximum chain capacity

- Question: if *r* is the throughput when node 1 transmits alone towards node 2, what is the maximum packet delivery rate between 1 and 8, assuming ideal packet scheduling?
- Answer: if tx order is 1-2-3-4-5-6-7-8, we have 7 tx before a packet delivery -> max thr = r/7
- Answer: we can exploit simultaneous tx! After a transient tx order cyclically is (1,4,7)-(2,5)-(3,6): we have 3 tx before a packet delivery -> max thr = r/3

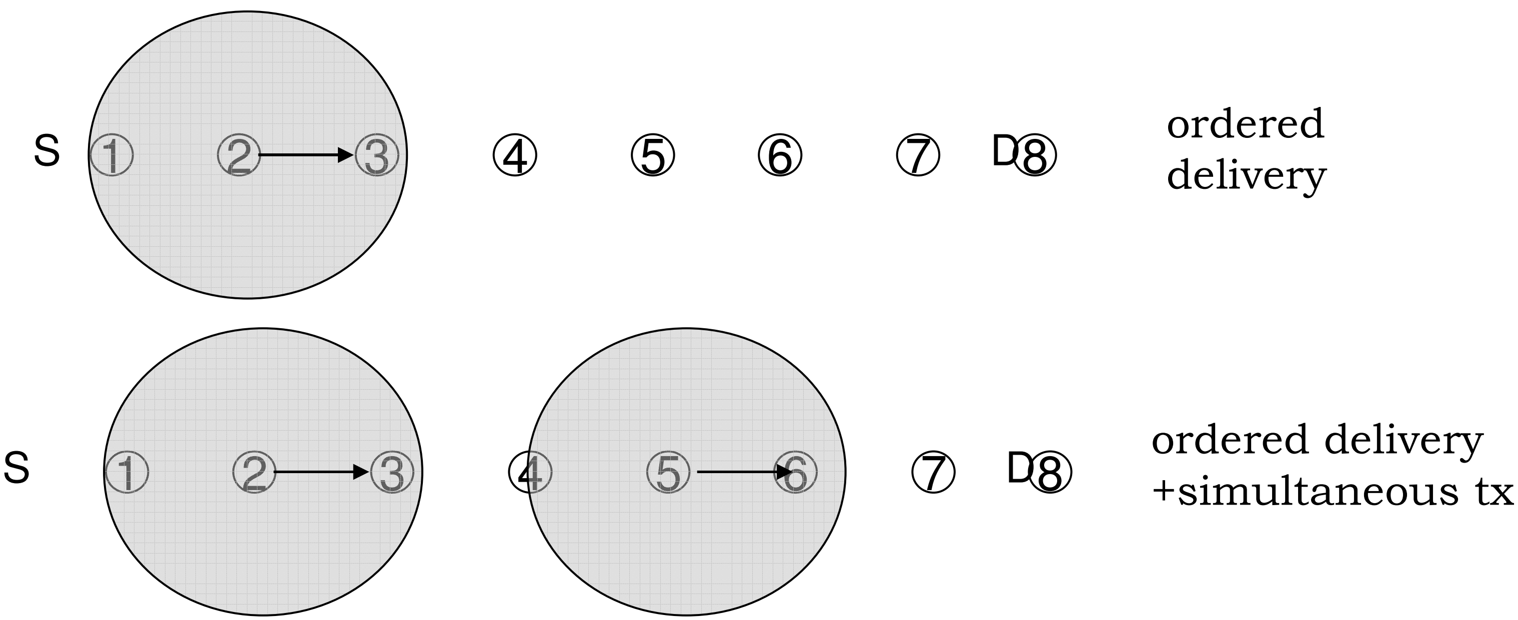

Maximum chain capacity

- Question: if *r* is the throughput when node 1 transmits alone towards node 2, what is the maximum packet delivery rate between 1 and 8, assuming ideal packet scheduling?
- Answer: if tx order is 1-2-3-4-5-6-7-8, we have 7 tx before a packet delivery -> max thr = r/7
- Answer: we can exploit simultaneous tx! After a transient tx order cyclically is (1,4,7)-(2,5)-(3,6): we have 3 tx before a packet delivery -> max thr = r/3

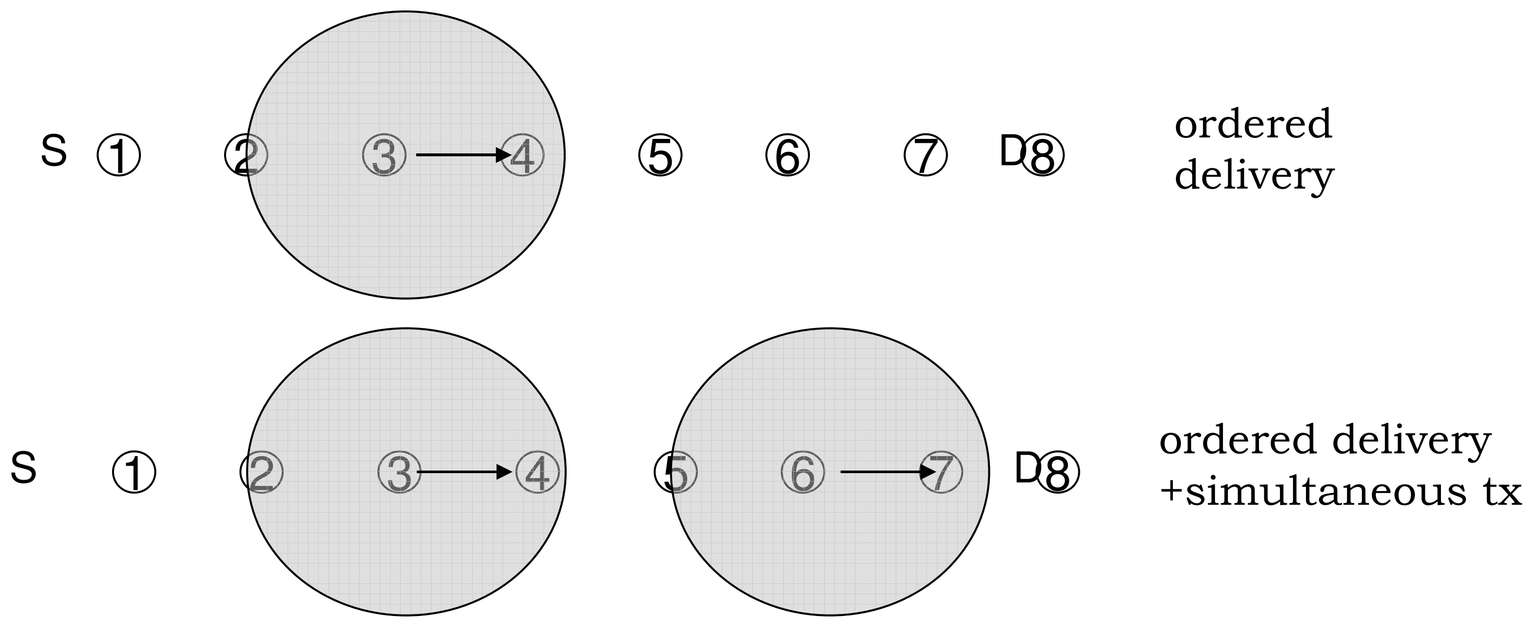

Maximum chain capacity

- Question: if *r* is the throughput when node 1 transmits alone towards node 2, what is the maximum packet delivery rate between 1 and 8, assuming ideal packet scheduling?
- Answer: if tx order is 1-2-3-4-5-6-7-8, we have 7 tx before a packet delivery -> max thr = r/7
- Answer: we can exploit simultaneous tx! After a transient tx order cyclically is (1,4,7)-(2,5)-(3,6): we have 3 tx before a packet delivery -> max thr = r/3

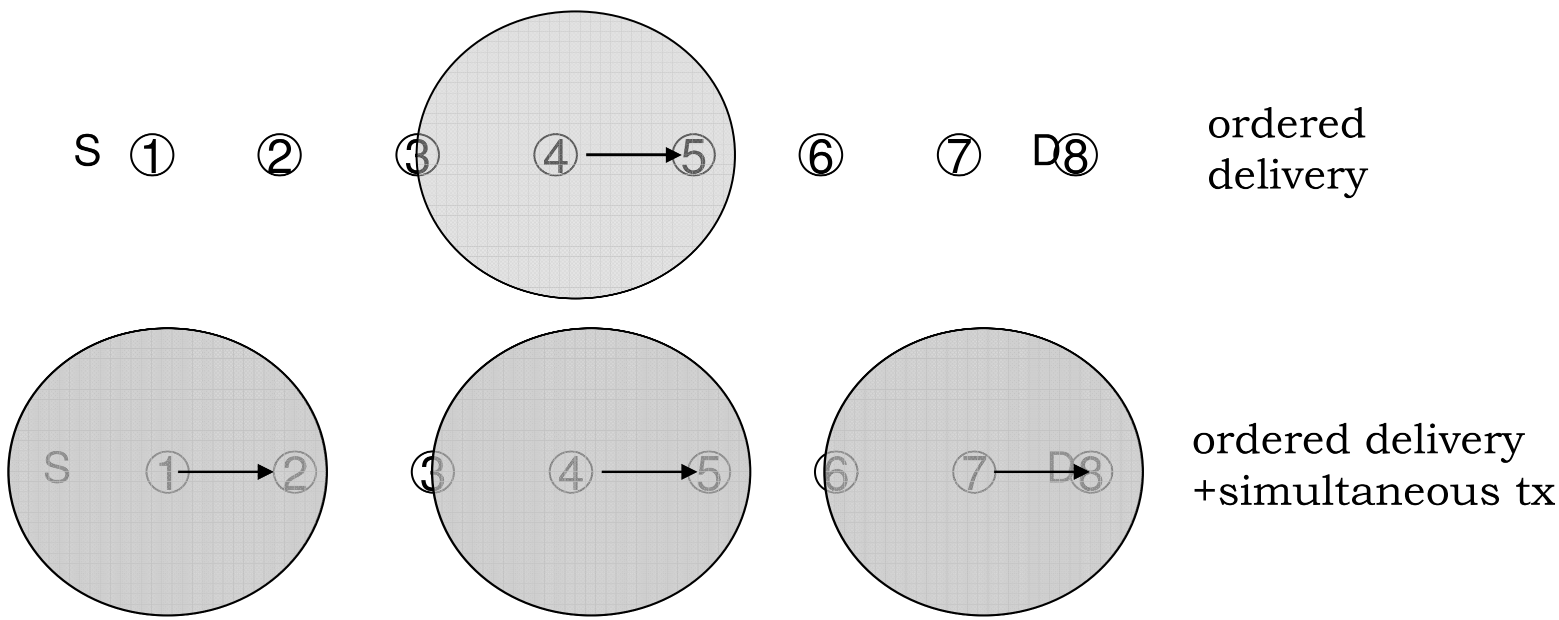

Maximum chain capacity

- Question: if *r* is the throughput when node 1 transmits alone towards node 2, what is the maximum packet delivery rate between 1 and 8, assuming ideal packet scheduling?
- Answer: if tx order is 1-2-3-4-5-6-7-8, we have 7 tx before a packet delivery -> max thr = r/7
- Answer: we can exploit simultaneous tx! After a transient tx order cyclically is (1,4,7)-(2,5)-(3,6): we have 3 tx before a packet delivery -> max thr = r/3

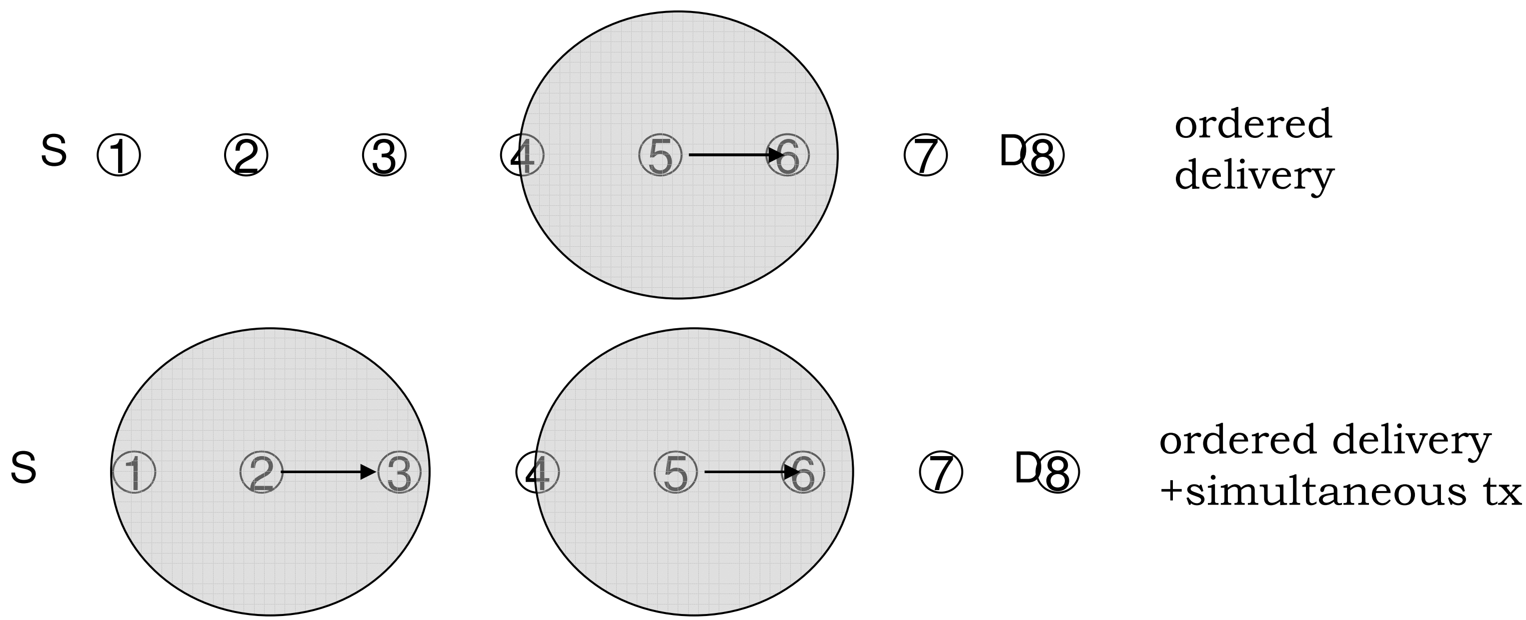

Maximum chain capacity

- Question: if *r* is the throughput when node 1 transmits alone towards node 2, what is the maximum packet delivery rate between 1 and 8, assuming ideal packet scheduling?
- Answer: if tx order is 1-2-3-4-5-6-7-8, we have 7 tx before a packet delivery -> max thr = r/7
- Answer: we can exploit simultaneous tx! After a transient tx order cyclically is (1,4,7)-(2,5)-(3,6): we have 3 tx before a packet delivery -> max thr = r/3

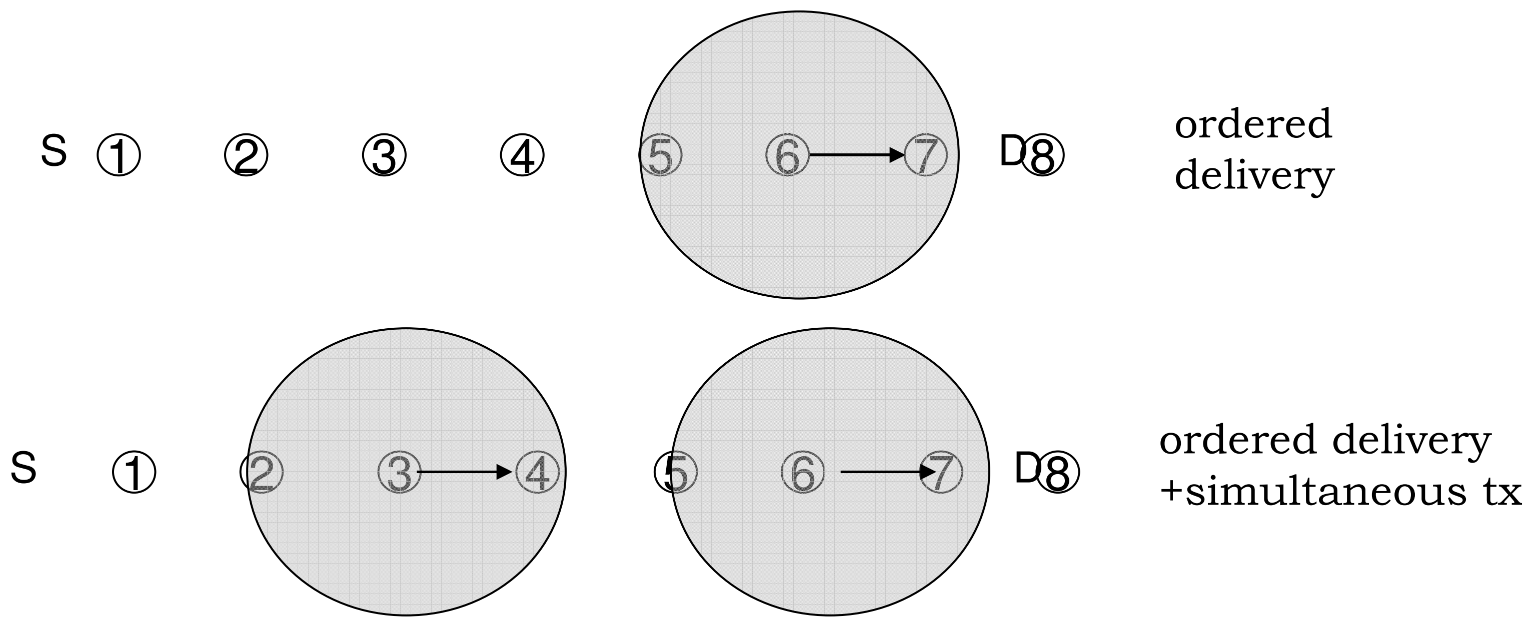

Maximum chain capacity

- Question: if *r* is the throughput when node 1 transmits alone towards node 2, what is the maximum packet delivery rate between 1 and 8, assuming ideal packet scheduling?
- Answer: if tx order is 1-2-3-4-5-6-7-8, we have 7 tx before a packet delivery -> max thr = r/7
- Answer: we can exploit simultaneous tx! After a transient tx order cyclically is (1,4,7)-(2,5)-(3,6): we have 3 tx before a packet delivery -> max thr = r/3

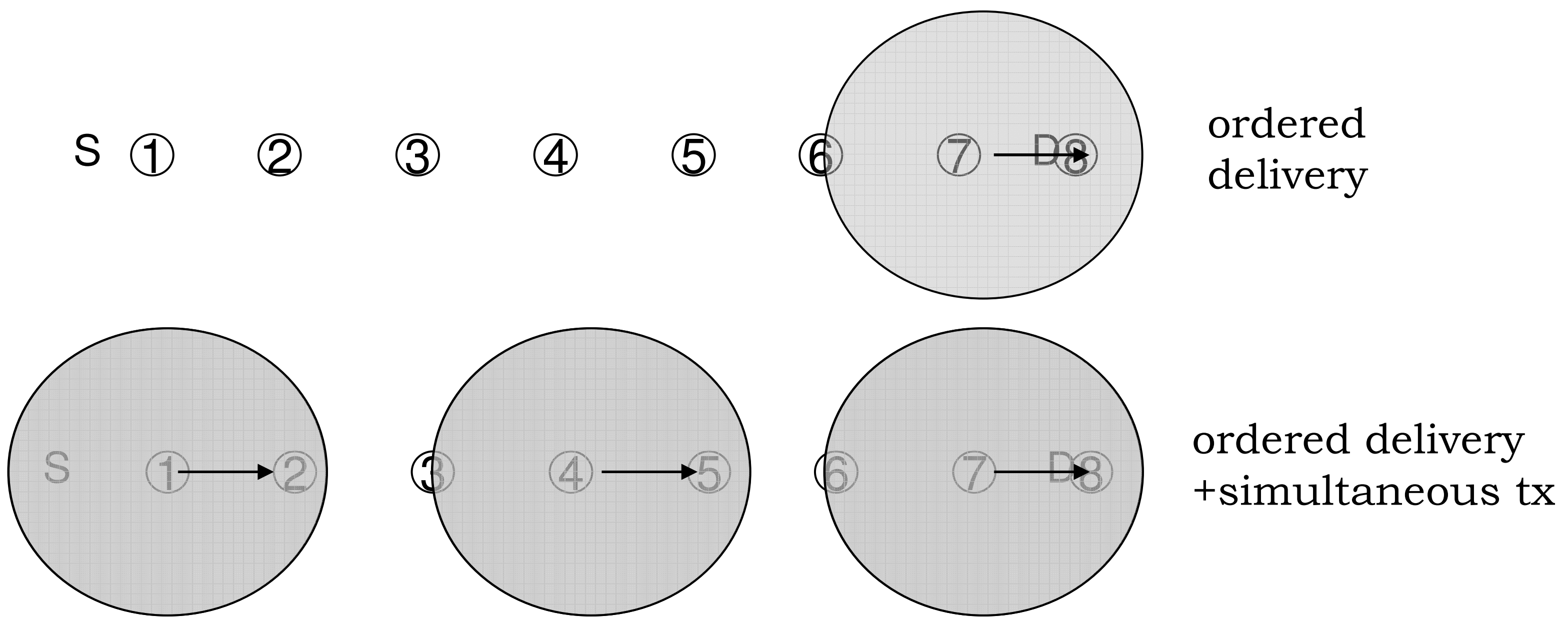

Actual chain capacity

- ➔ **DCF is totally distributed! No ideal scheduling among the node transmissions**
- ➔ **Dramatic hidden node problem, especially for the first nodes of the chain**
  - ⇨ Along the chain, is rare that contiguous nodes are simultaneously active
  - ⇨ Collisions on the back of the packet flow direction (e.g. collision @node 2, not @node 4!)

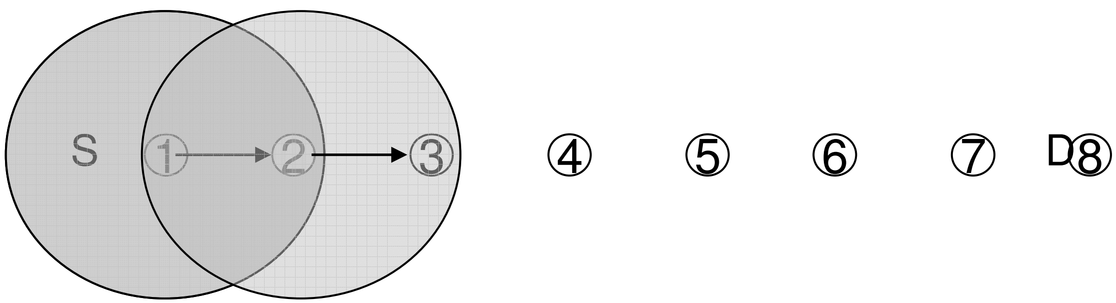

1. node 1 tx its first packet
2. node 1 and node 2 contend for the next channel access
3. After the first node 2 successful tx, it is very likely that next node 1 tx is originated during ongoing node 3 tx!

Actual chain capacity

- **DCF is totally distributed! No ideal scheduling among the node transmissions**
- **Dramatic hidden node problem, especially for the first nodes of the chain**
  - Along the chain, is rare that contiguous nodes are simultaneously active
  - Collisions on the back of the packet flow direction (e.g. collision @node 2, not @node 4!)

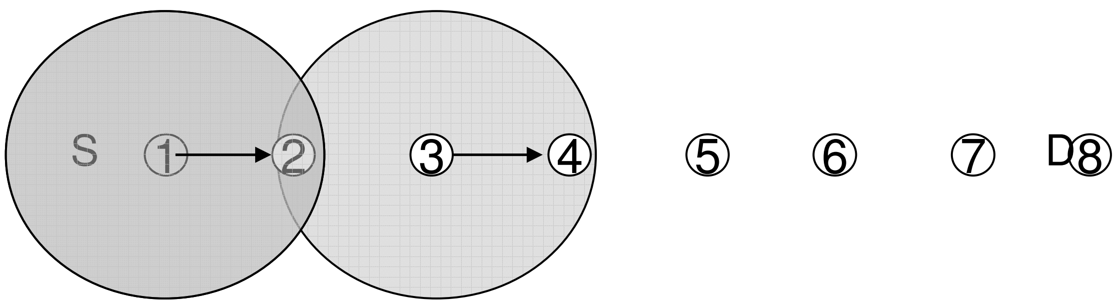

1. node 1 tx its first packet
2. node 1 and node 2 contend for the next channel access
3. After the first node 2 successful tx, it is very likely that next node 1 tx is originated during ongoing node 3 tx!

Can RTS/CTS help?

Two different collision events @node 2:

- ⇨ node 3 starts its tx during ongoing node 1 tx;
- ⇨ node 1 starts its tx during ongoing node 3 tx

➔ **RTS/CTS do not solve the second collision event, which is the most common!**

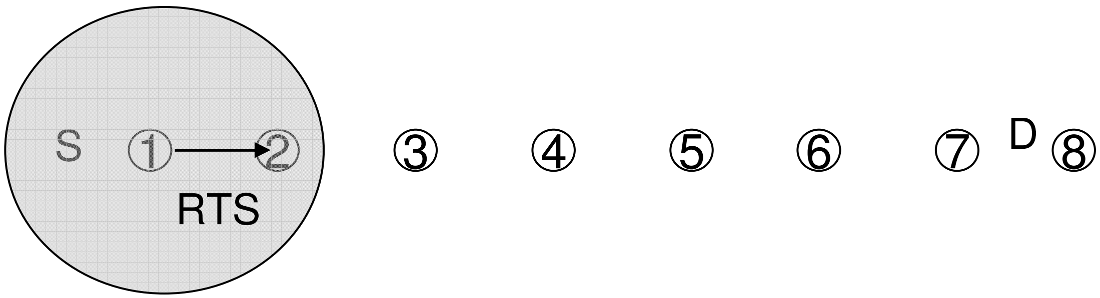

1. node 1 tx RTS
2. node 2 replies with a CTS packet which blocks node 3 tx
3. node 1 tx DATA: Ok!

1. node 3 tx RTS
2. node 4 replies with a CTS packet blocking node 5 tx
3. node 1 tx RTS: collisions at node 2! (often more subsequent RTS collisions during the same node 3 data tx)

Can RTS/CTS help?

Two different collision events @node 2:

- ⇨ node 3 starts its tx during ongoing node 1 tx;
- ⇨ node 1 starts its tx during ongoing node 3 tx

➔ **RTS/CTS do not solve the second collision event, which is the most common!**

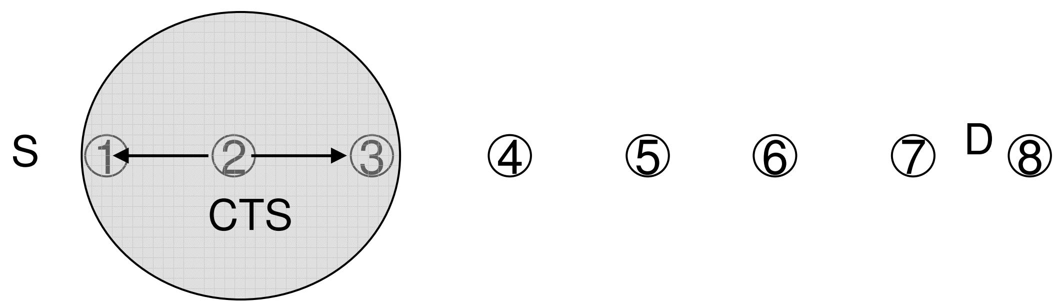

1. node 1 tx RTS
2. node 2 replies with a CTS packet which blocks node 3 tx
3. node 1 tx DATA: Ok!

1. node 3 tx RTS
2. node 4 replies with a CTS packet blocking node 5 tx
3. node 1 tx RTS: collisions at node 2! (often more subsequent RTS collisions during the same node 3 data tx)

Can RTS/CTS help?

Two different collision events @node 2:

- ⇨ node 3 starts its tx during ongoing node 1 tx;
- ⇨ node 1 starts its tx during ongoing node 3 tx

➔ RTS/CTS do not solve the second collision event, which is the most common!

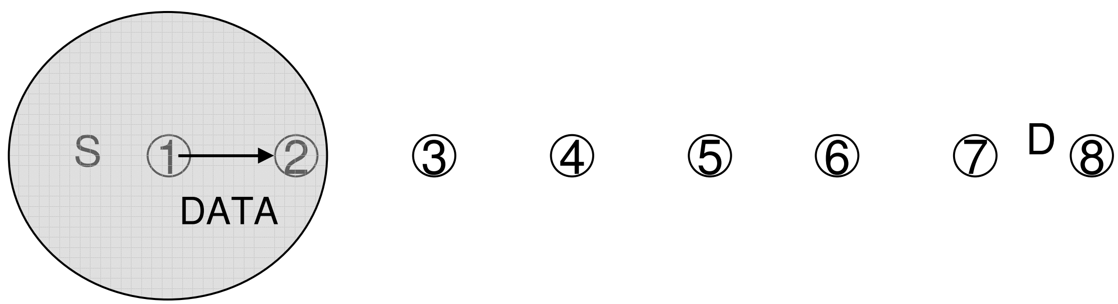

1. node 1 tx RTS
2. node 2 replies with a CTS packet which blocks node 3 tx
3. node 1 tx DATA: Ok!

1. node 3 tx RTS
2. node 4 replies with a CTS packet blocking node 5 tx
3. node 1 tx RTS: collisions at node 2! (often more subsequent RTS collisions during the same node 3 data tx)

Can RTS/CTS help?

Two different collision events @node 2:

- ⇨ node 3 starts its tx during ongoing node 1 tx;
- ⇨ node 1 starts its tx during ongoing node 3 tx

➔ **RTS/CTS do not solve the second collision event, which is the most common!**

1. node 1 tx RTS
2. node 2 replies with a CTS packet which blocks node 3 tx
3. node 1 tx DATA: Ok!

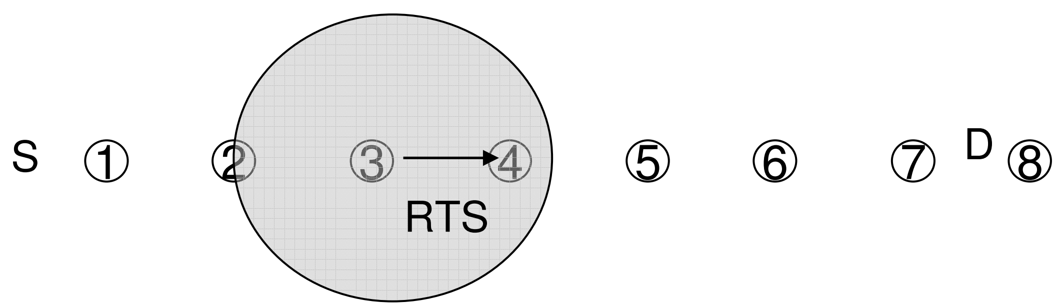

1. node 3 tx RTS
2. node 4 replies with a CTS packet blocking node 5 tx
3. node 1 tx RTS: collisions at node 2! (often more subsequent RTS collisions during the same node 3 data tx)

Can RTS/CTS help?

Two different collision events @node 2:

- ⇨ node 3 starts its tx during ongoing node 1 tx;
- ⇨ node 1 starts its tx during ongoing node 3 tx

➔ **RTS/CTS do not solve the second collision event, which is the most common!**

1. node 1 tx RTS
2. node 2 replies with a CTS packet which blocks node 3 tx
3. node 1 tx DATA: Ok!

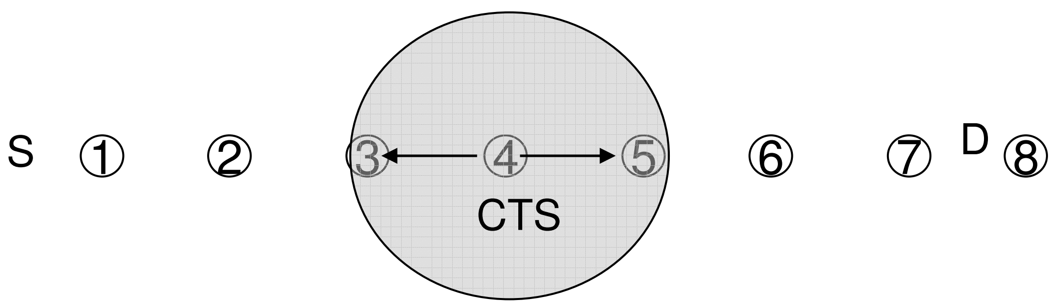

1. node 3 tx RTS
2. node 4 replies with a CTS packet blocking node 5 tx
3. node 1 tx RTS: collisions at node 2! (often more subsequent RTS collisions during the same node 3 data tx)

Can RTS/CTS help?

Two different collision events @node 2:

- ⇨ node 3 starts its tx during ongoing node 1 tx;
- ⇨ node 1 starts its tx during ongoing node 3 tx

➔ **RTS/CTS do not solve the second collision event, which is the most common!**

1. node 1 tx RTS
2. node 2 replies with a CTS packet which blocks node 3 tx
3. node 1 tx DATA: Ok!

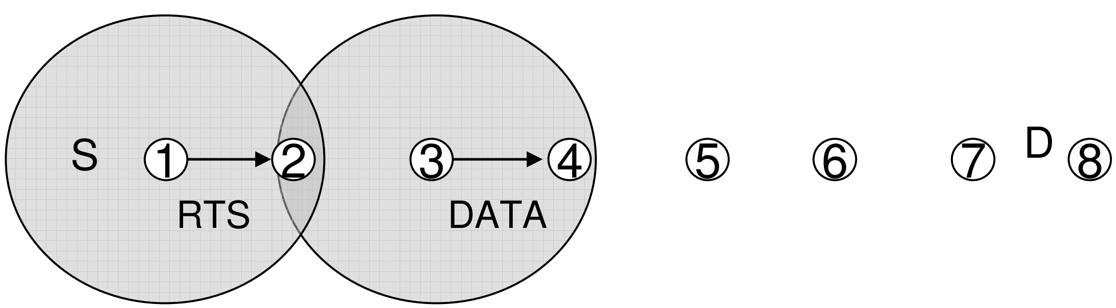

1. node 3 tx RTS
2. node 4 replies with a CTS packet blocking node 5 tx
3. node 1 tx RTS: collisions at node 2! (often more subsequent RTS collisions during the same node 3 data tx)

RTS/CTS Collision Times

in node chains

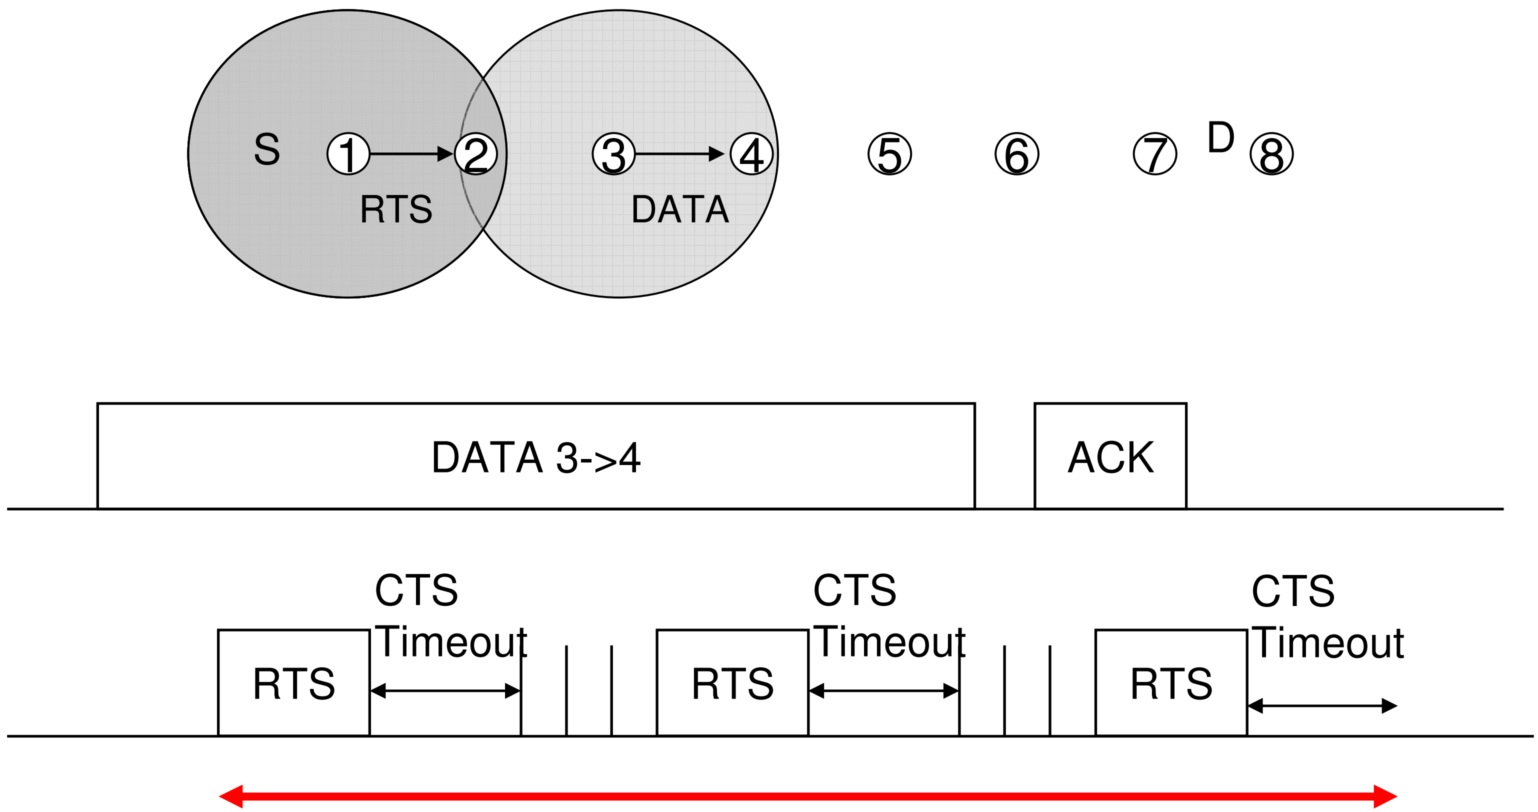

2 drawbacks:
Actual collision times are not reduced!
Because of multiple collisions, higher CW and higher next access delays!

Spatial reuse via directional antennas

- ➔ **Smart antennas/ switched beam may be effectively deployed over multiple transceiver APs**
  - ⇨ Possible capable of independent simultaneous TX/RX on all beams
- ➔ **Goal: enable simultaneous tx/rx in different beasm**
  - ⇨ Space-Division Multiple Access (SDMA)
- ➔ **Design Constraint: omnidirectional antennas on STA**
- ➔ **Not a problem: beam forming done at the AP (valid for both TX and RX directions)**

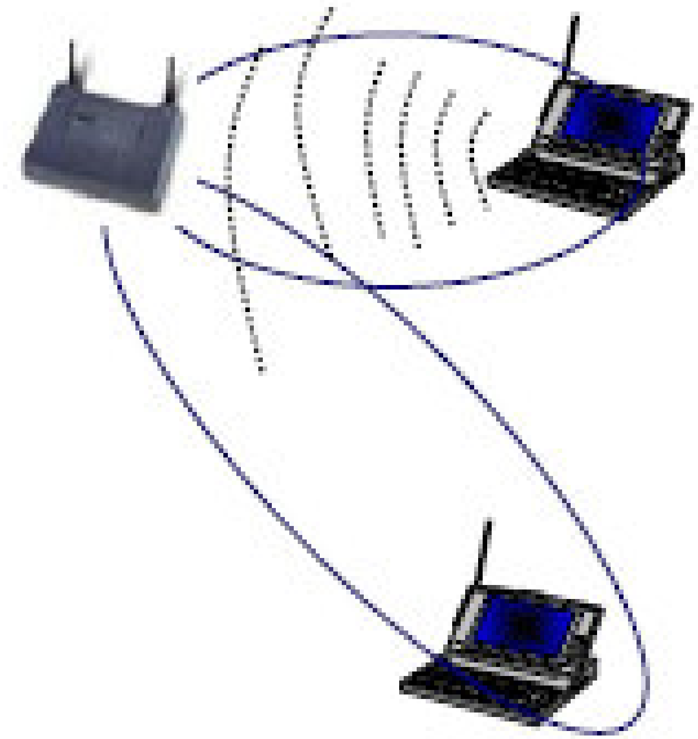

Cell Capacity

- If we complicate the AP structure, with multi transceivers and directional antennas, we can multiply the radio resources available in a given cell
- Omni-directional vs. Directive Beams: more beams, more capacity!
  - Does it work with standard DCF??

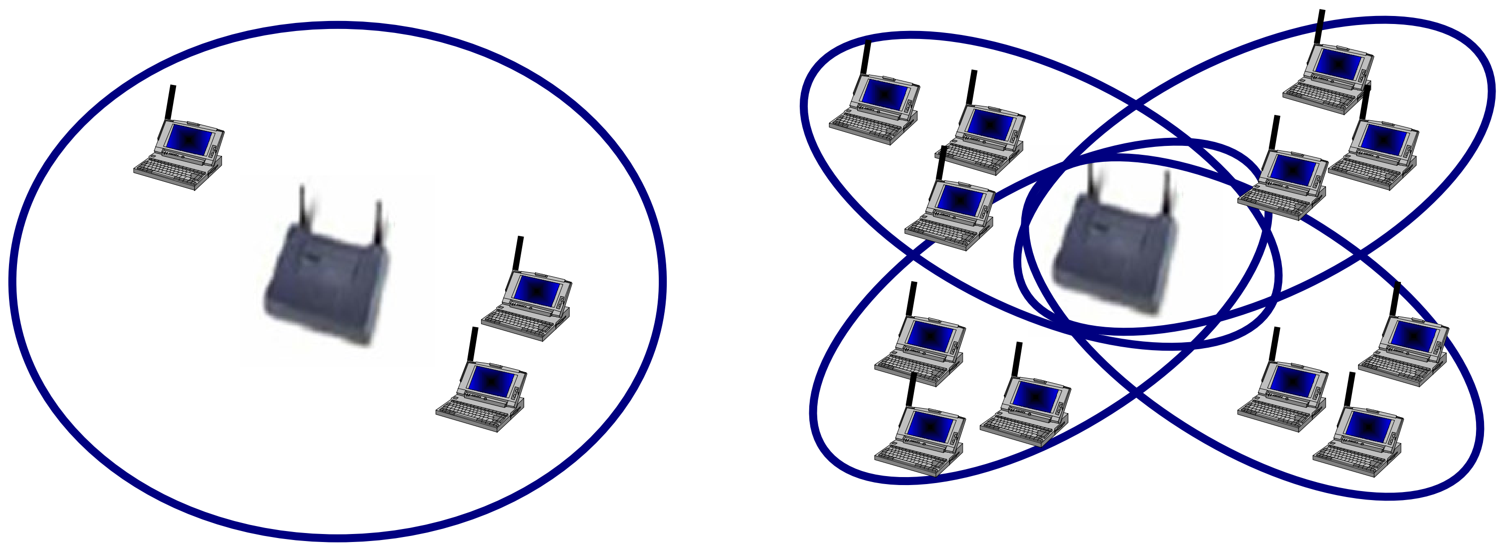

Actual scenario: some thoughts

- **How much directive antennas may increase the capacity of a cell?**
  - We are not interested here to increase covered distance
- **Working assumption**
  - 1 central AP;
  - Ideal operation of directional antennas
  - Many STA, all in reciprocal visibility
    - Antenna technology used to increase capacity; no power control issues considered
  - Assume STA positions known

Simultaneous uplink/downlink TX

1. **AP is transmitting to STA A**
2. **STA B performs carrier sensing**
3. **STA B sends omnidirectional DATA**
4. **STA B DATA destroys STA A ongoing reception**

➔ **Conclusion**
- ⇨ We need to prevent TX from B
  - → E.g. via omnidirectional CTS from A
- ⇨ If all STAs are in range, simultaneous uplink/downlink TXs impossible

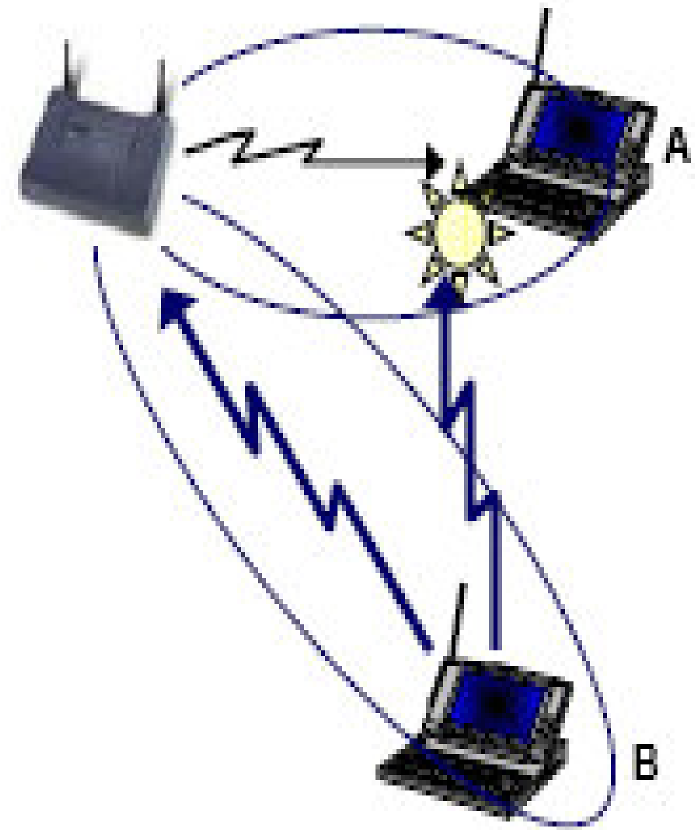

Simultaneous downlink

1. AP is transmitting to STA A and STA B simultaneously
2. DATA to B ends; after a SIFS B sends ACK
3. .. Which destroys A reception

➔ **Conclusion**

⇨ Unless accurate scheduling considered, simultaneous downlink TX are not possible

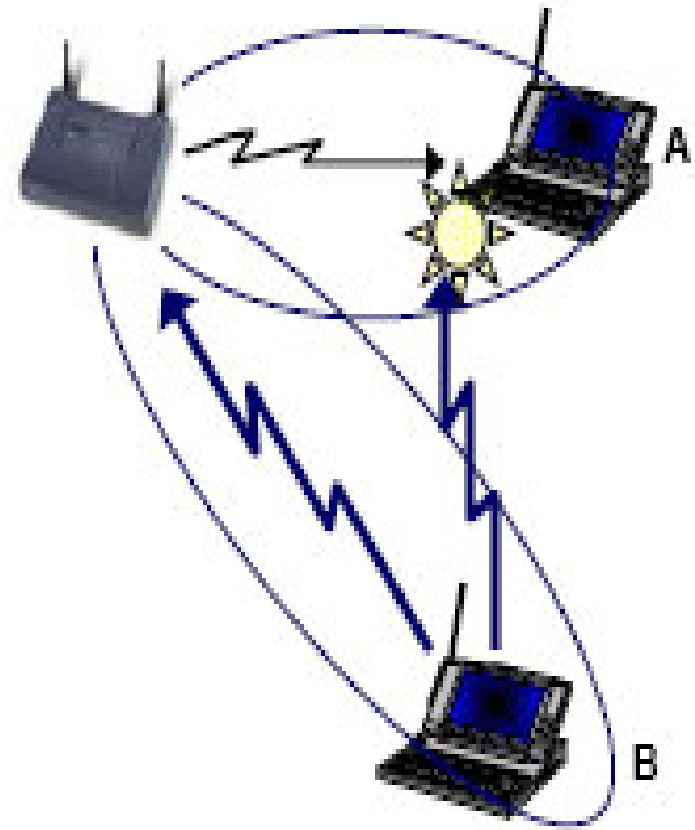

Simultaneous uplink

1. **B transmits to AP**
2. **A might transmit to AP too..**
   Note that subsequent ACK would be directed and would not interfere
3. **..but senses the channel busy**

➔ **Conclusion**
   - ⇨ Exposed terminal problem magnified
   - ⇨ Simultaneous uplink transmissions are not possible

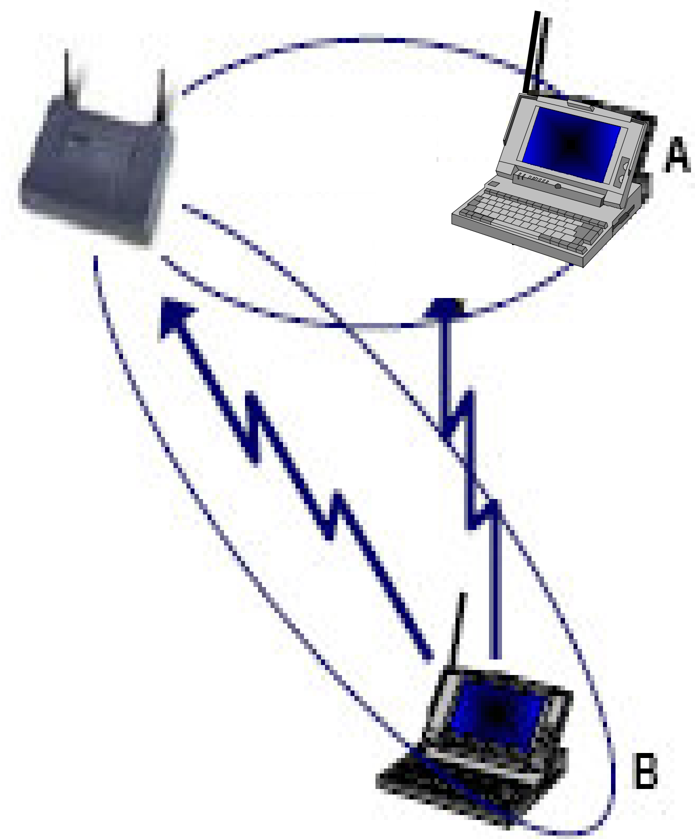

Summarizing…

- **The asynchronoud DCF handshake is way far from being suited to support SDMA**
  - We have just proven that, in full coverage, only a SINGLE transmission at a time my occur into a cell
- **Solutions:**
  - Centralized MAC;
  - Power control
  - New MAC (Throw DCF away!)

Multiple Radio MAC

- ➔ **Taking dinamicity in the MAC: multi-channel MAC**
  - →**[Nasipuri, Zhuang, Das, 1999]; [Jain, Das, Nasipuri, 2001]**
  - →**[Tseng, Wu, Lin, 2001]**
  - →**[Hung, Law, Leon-Garcia, 2002]**

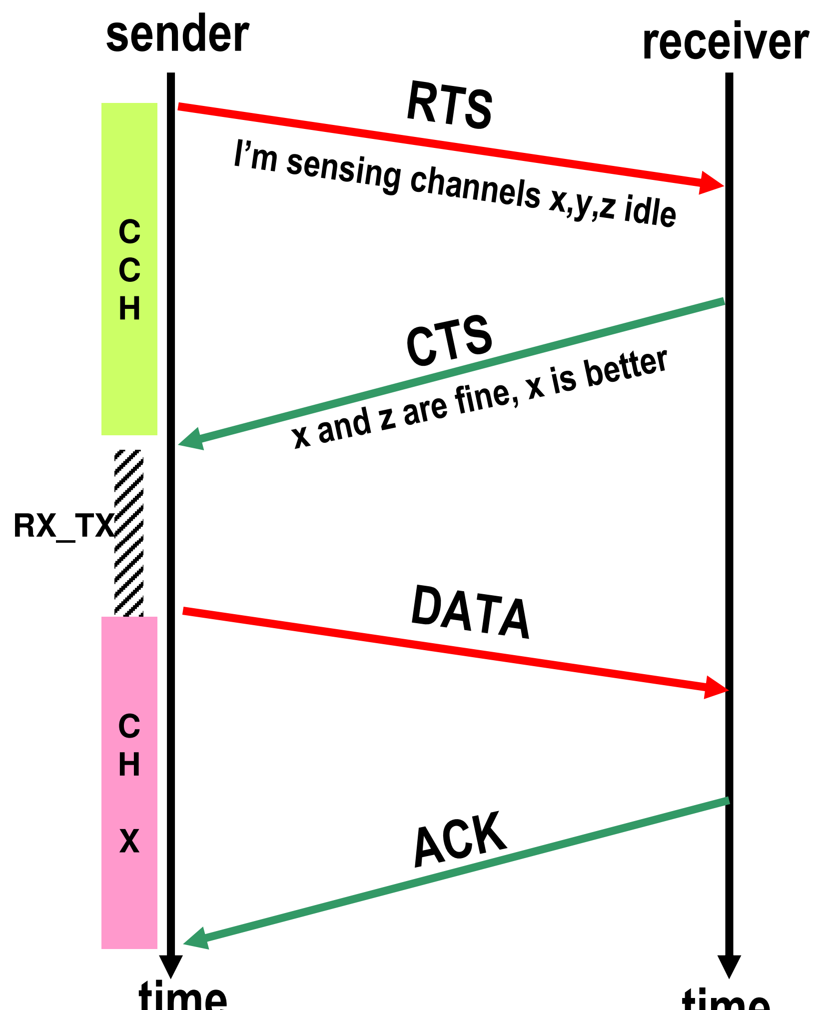

- ➔ **Multiple channels available**
- ➔ **DATA transmitted on channel selected via (modified) RTS/CTS handshake**
  - ⇨ RTS/CTS handshake on Common Control (signalling) Channel

Implementation issues

- ➔ **Implementation transparent to MAC**
  - →Multichannel handshake coded into PLCP header
    - » [Technical report in italian project FIRB-PRIMO]
  - →MAC sees a unique channel

- ➔ **Technical issues (not discussed in papers)**
  - ⇨ Multi-channel carrier sense
    - →Hard with commercial components…
  - ⇨ Timing constraints for channel switching
    - →Again, many products do not support required timing

Multi Channel MAC

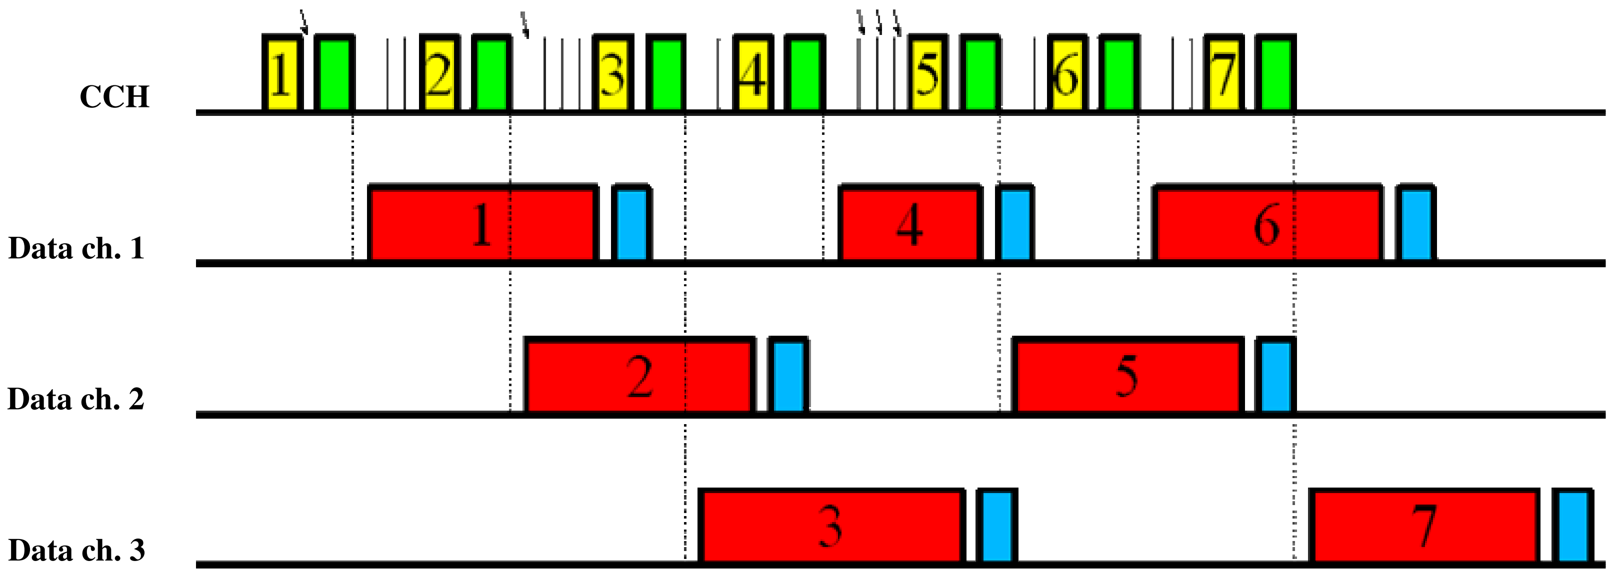

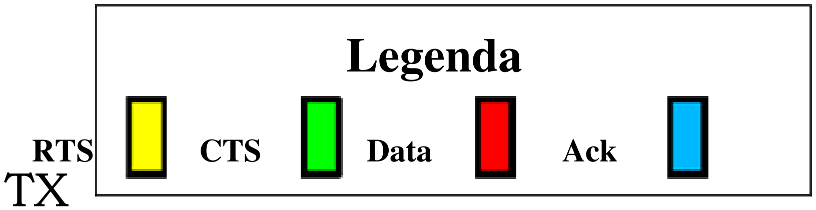

- ⇨ Legacy RTS/CTS handshake
  - →On control channel, only
- ⇨ Limited exploitation of parallel TX
  - →Approach not exploited to its full capabilities
  - →Channel separation wastes capacity
- ⇨ Tradeoffs required
  - →How much bandwidth to (bottleneck) signalling channel?

Rate optimization

Control channel data rate cannot be arbitrarily low, in order to avoid data channel wastes

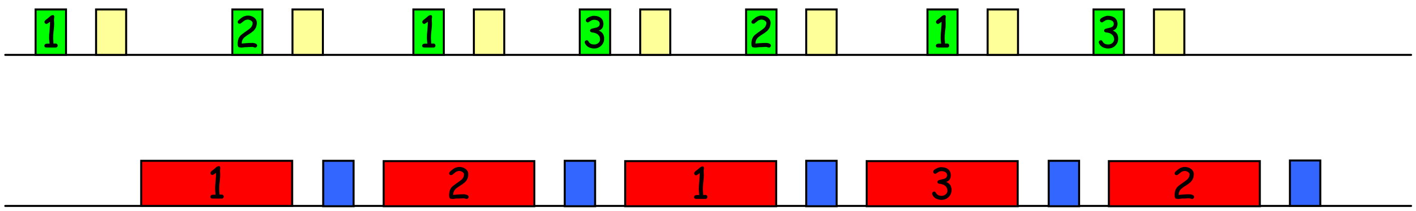

fully utilized data channel

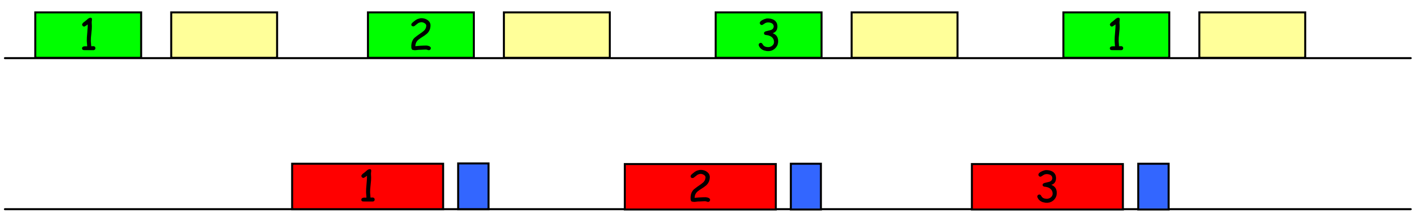

Resource wastes due to lack of reservations

QoS Support

802.11 MAC evolution

(802.11e, finalized in dicember 2005)

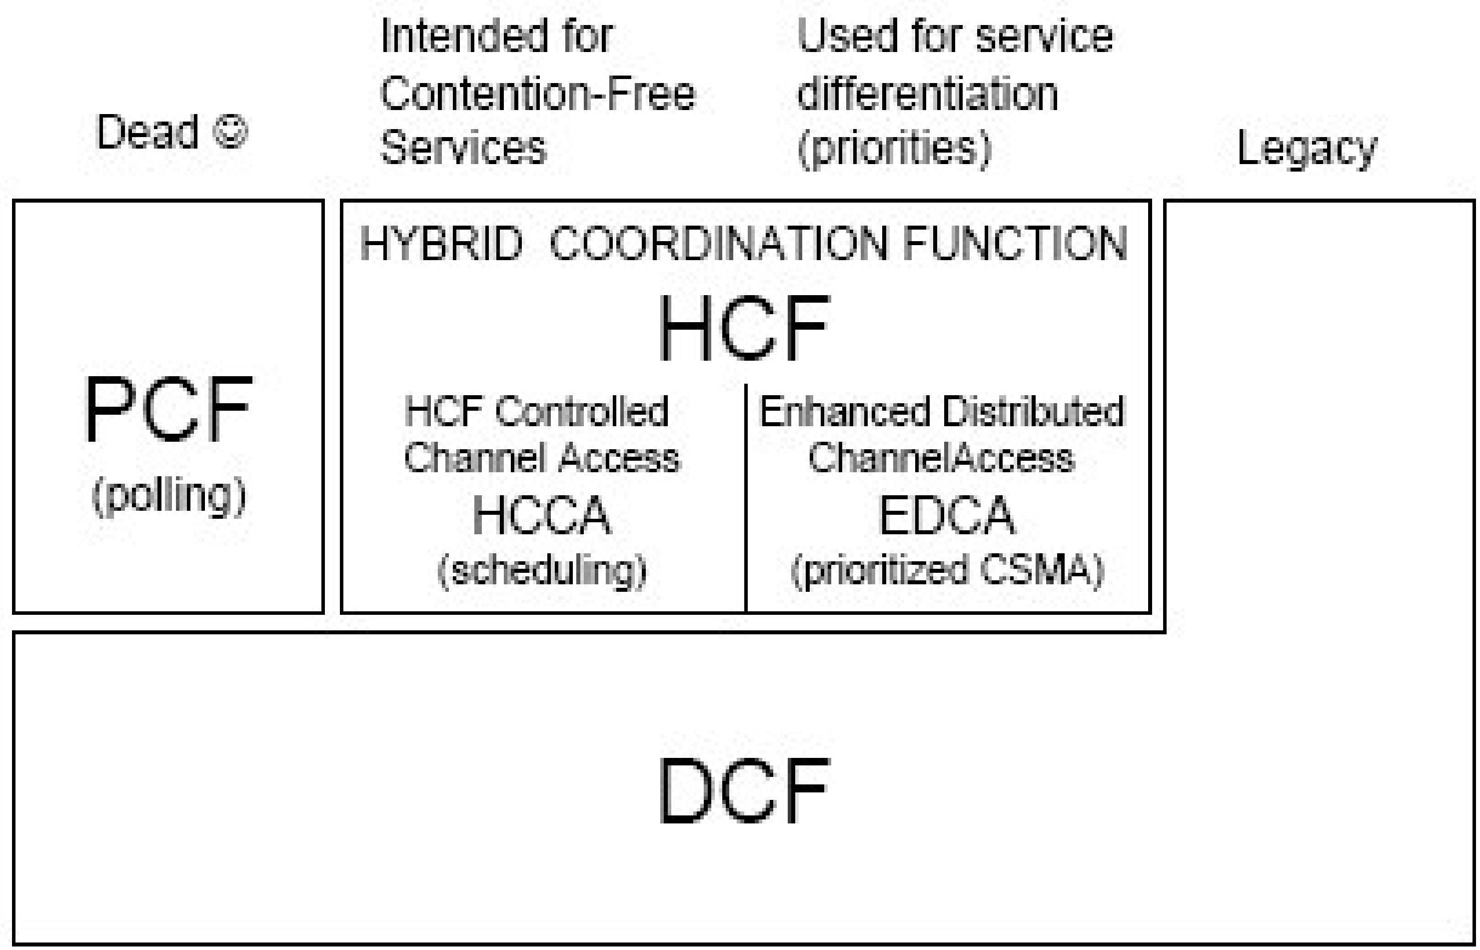

Multiple Queues

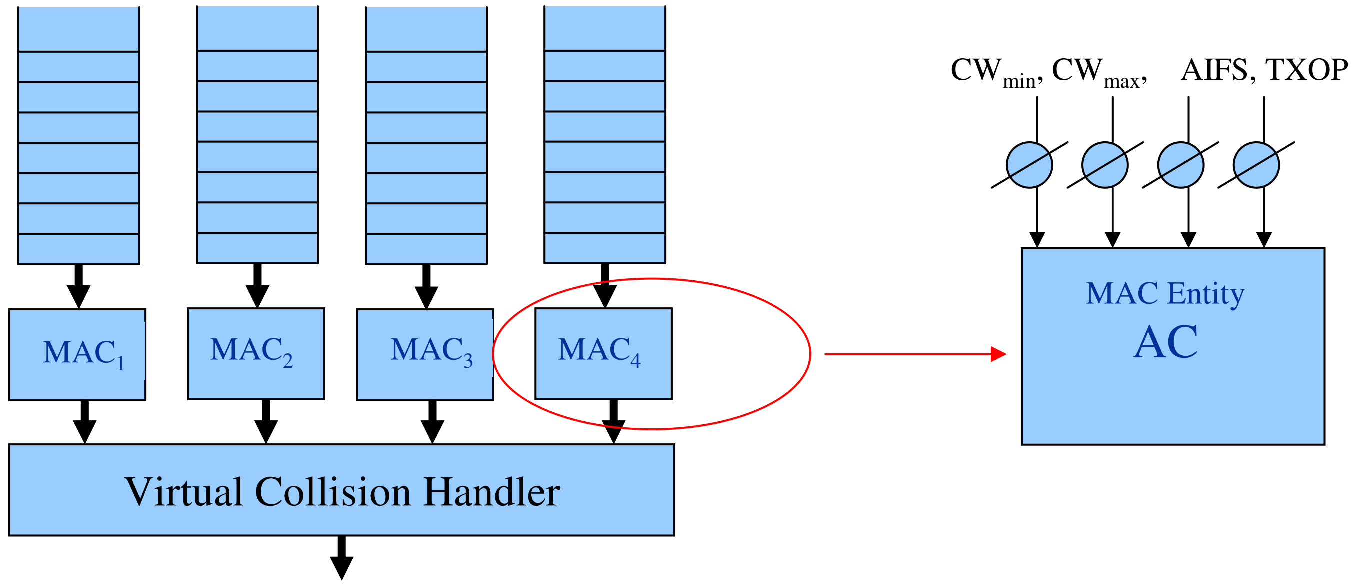

- ➔ **4 Access Categories**
  - ⇨ Mapping the 8 priority levels provided by 802.1p
  - ⇨ Different channel access probability through different access parameters
- ➔ **Independently operated as multiple MAC**
  - ⇨ Queues in the same station can (virtually) collide!

Distributed Prioritization: channel accesses

➔ **More channel accesses to High Priority stations reducing the backoff expiration times**

- ⇨ By giving probabilistically lower backoff counters (CWmin, CWmax)
- ⇨ By giving deterministically lower backoff resume times (AIFS)

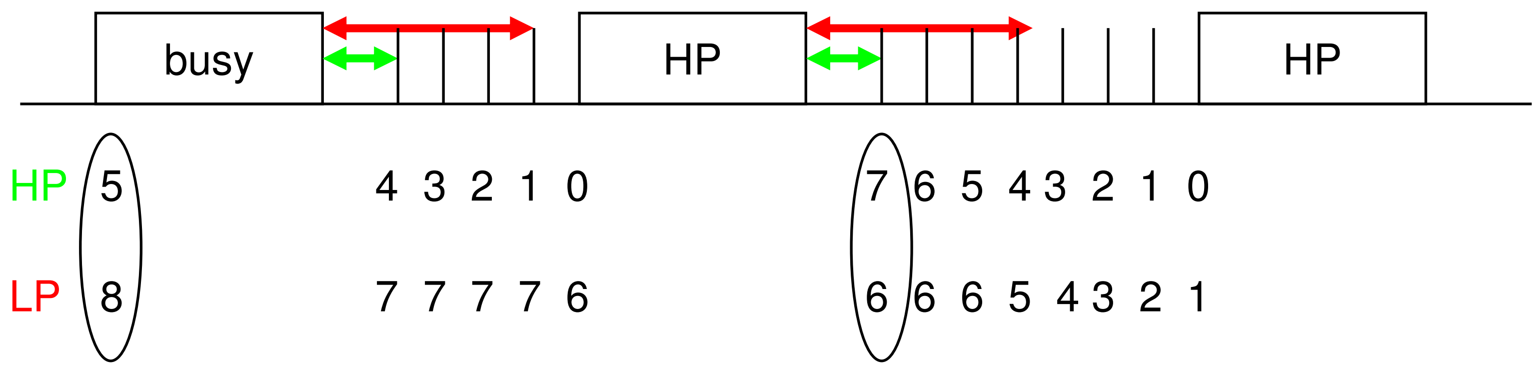

N.B. Tunable CWmin can also be used for performance optimizations as a function of the network load!!

Distributed Prioritization: transmission grants

- ➔ **Given the channel access probability, we can also differentiate the number of packet transmissions allowed for the stations which wins the contention**
- ➔ **More transmissions opportunities back-to-back to High Priority stations**
  - ⇨ Channel grants not on MSDU basis, but in terms of “channel holding times”

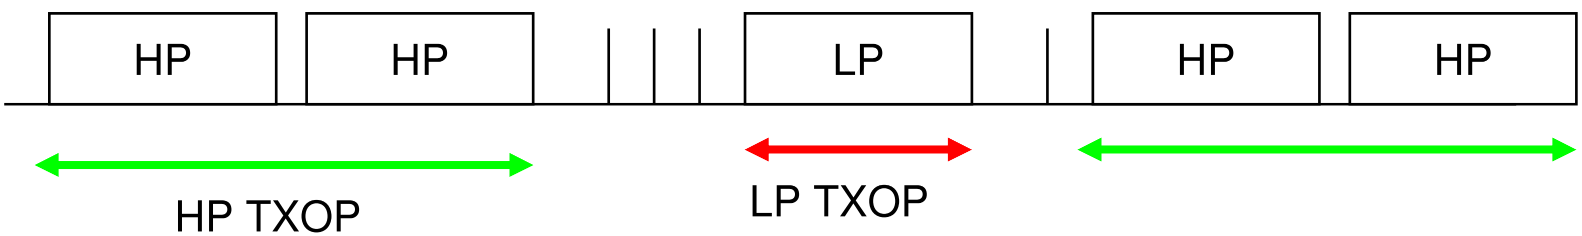

TXOP not only for throughput repartition, but also for efficiency improvements!

802.11: Old MAC and New PHYs..

➔ **In standard DCF, channel accesses are *packet oriented*: each MSDU transmission requires a different access**

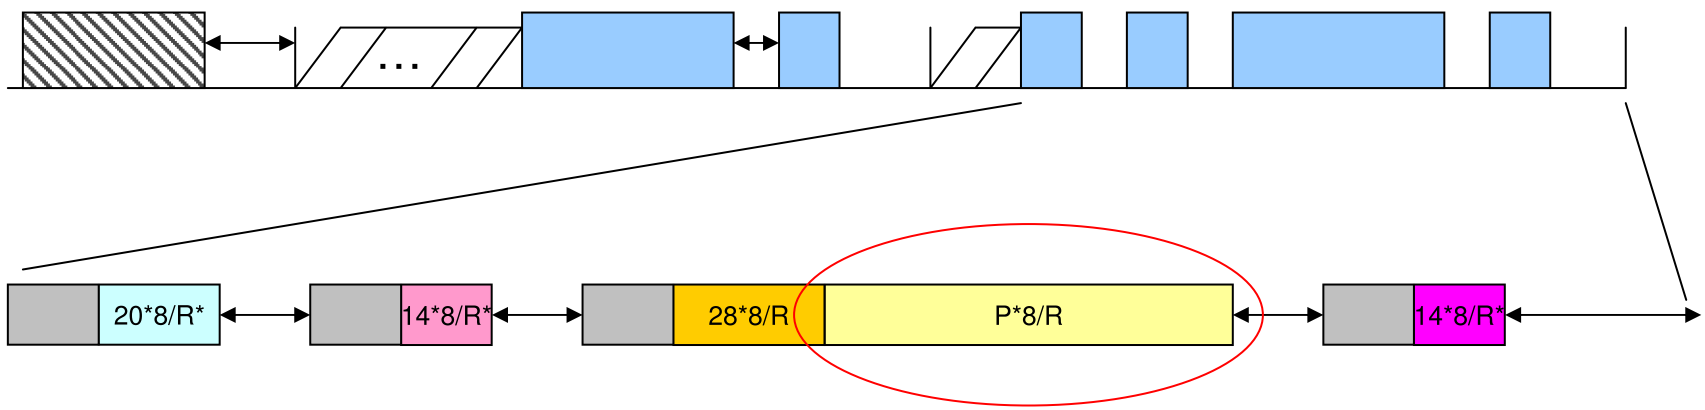

- Channel wastes are due to both PHY layer constraints and MAC operations:
  - SIFS, DIFS, SlotTime, Preamble, TX rates R and R*
  - RTS, CTS, ACK, # of bk slots, Collision Probability
- New PHYs allow higher TX rates..
  - Overheads are not reduced proportionally

Overheads @ different rates

(Packet=1500 bytes)

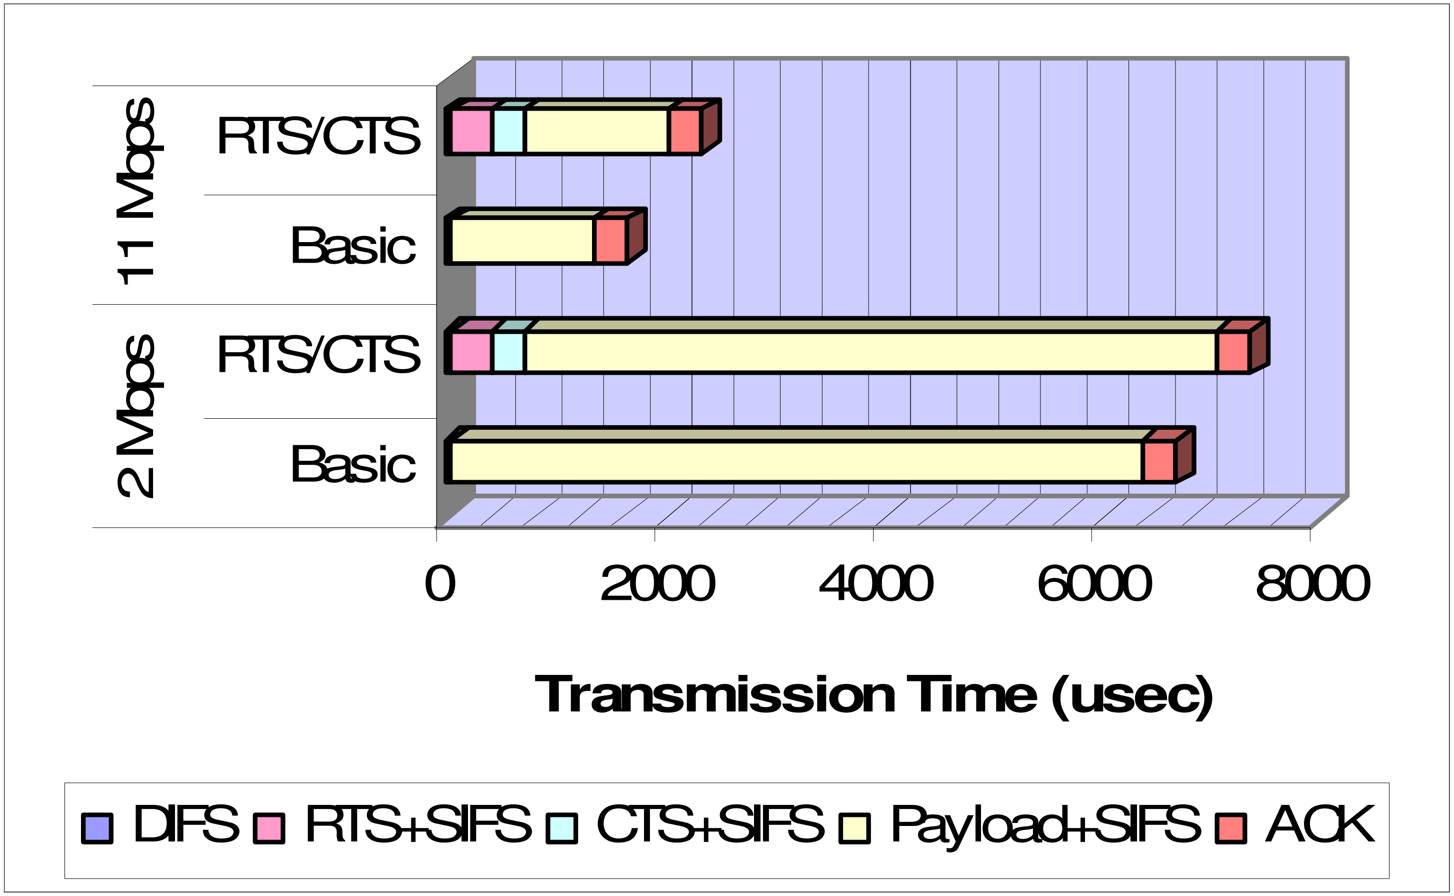

System efficiency degrades for high data rates!

802.11e transmission extensions

- **Key idea: the system efficiency improves by maximizing the payload transmission in each channel access (since overheads are reduced proportionally reduced)**
  - But maximum payload size is limited to 2304 bytes!

- **TXOP & BACK:**
  - Perform multiple transmissions in burst in each channel access
  - Acknowledge more packet transmissions with a cumulative ACK

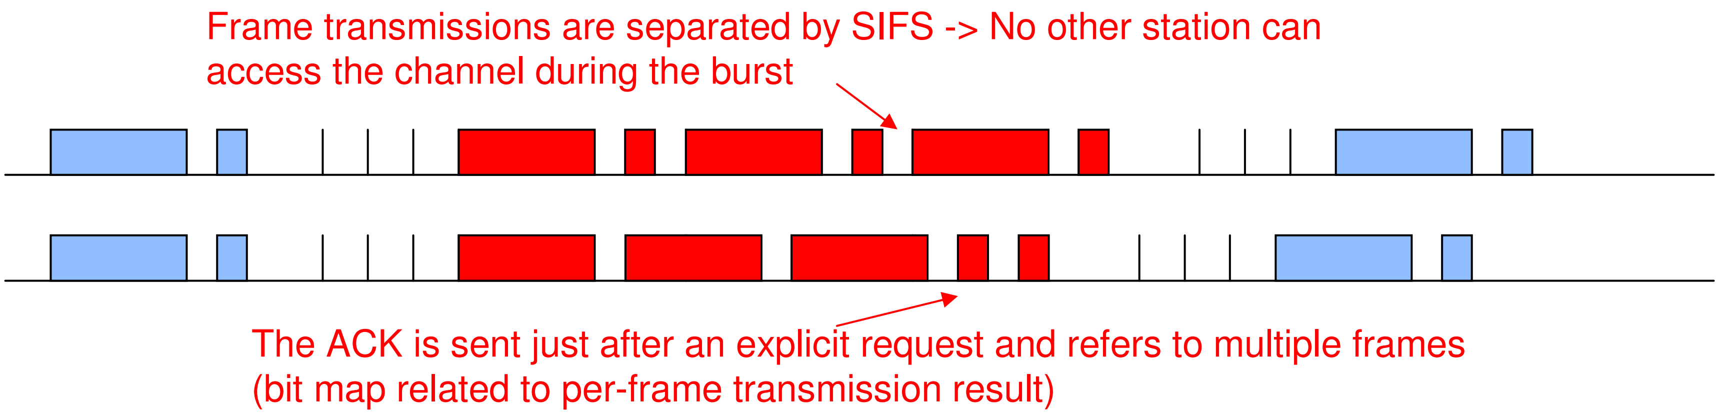

ACK Aggregation: does it work?

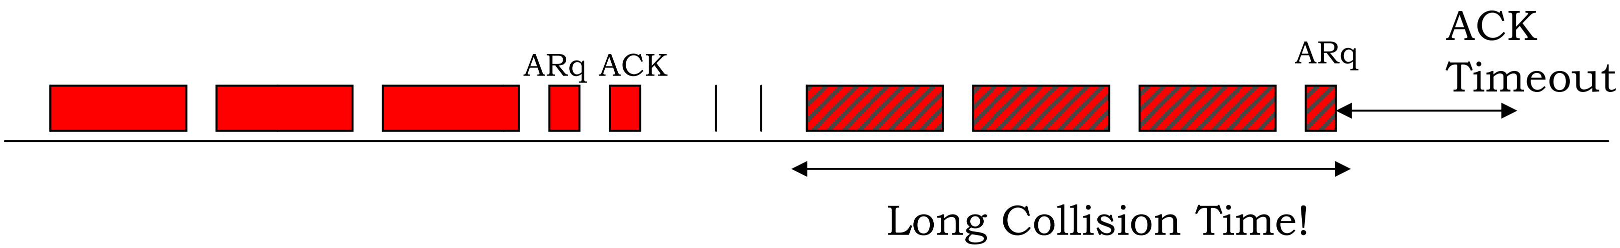

Collisions are revealed only after the transmission of the ACK Request (ARq) frame -> Collision times increase significantly.

Since only the Head Of Burst frame is subject to possible collisions, better strategies could be:

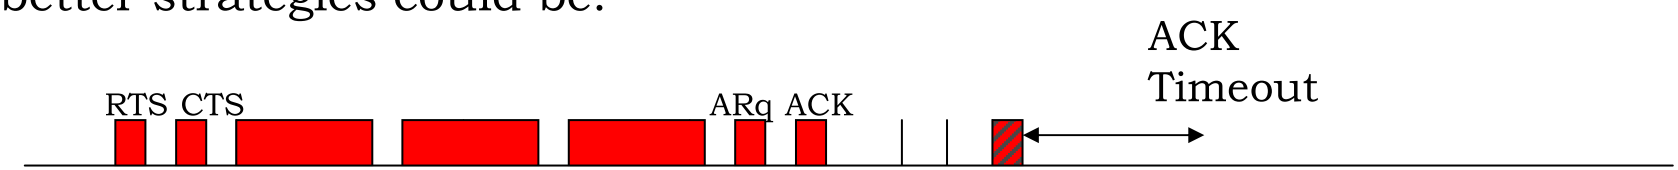

a) Preliminary RTS/CTS exchange in order to confirm the successful access

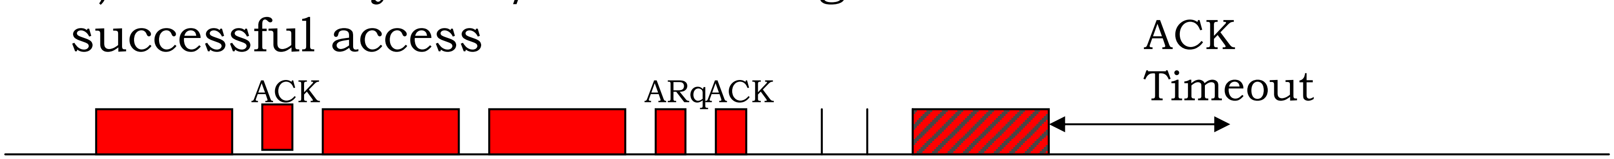

b) Explicit ACK for the first Data Frame before start the TX burst

Different Access and ACK policies

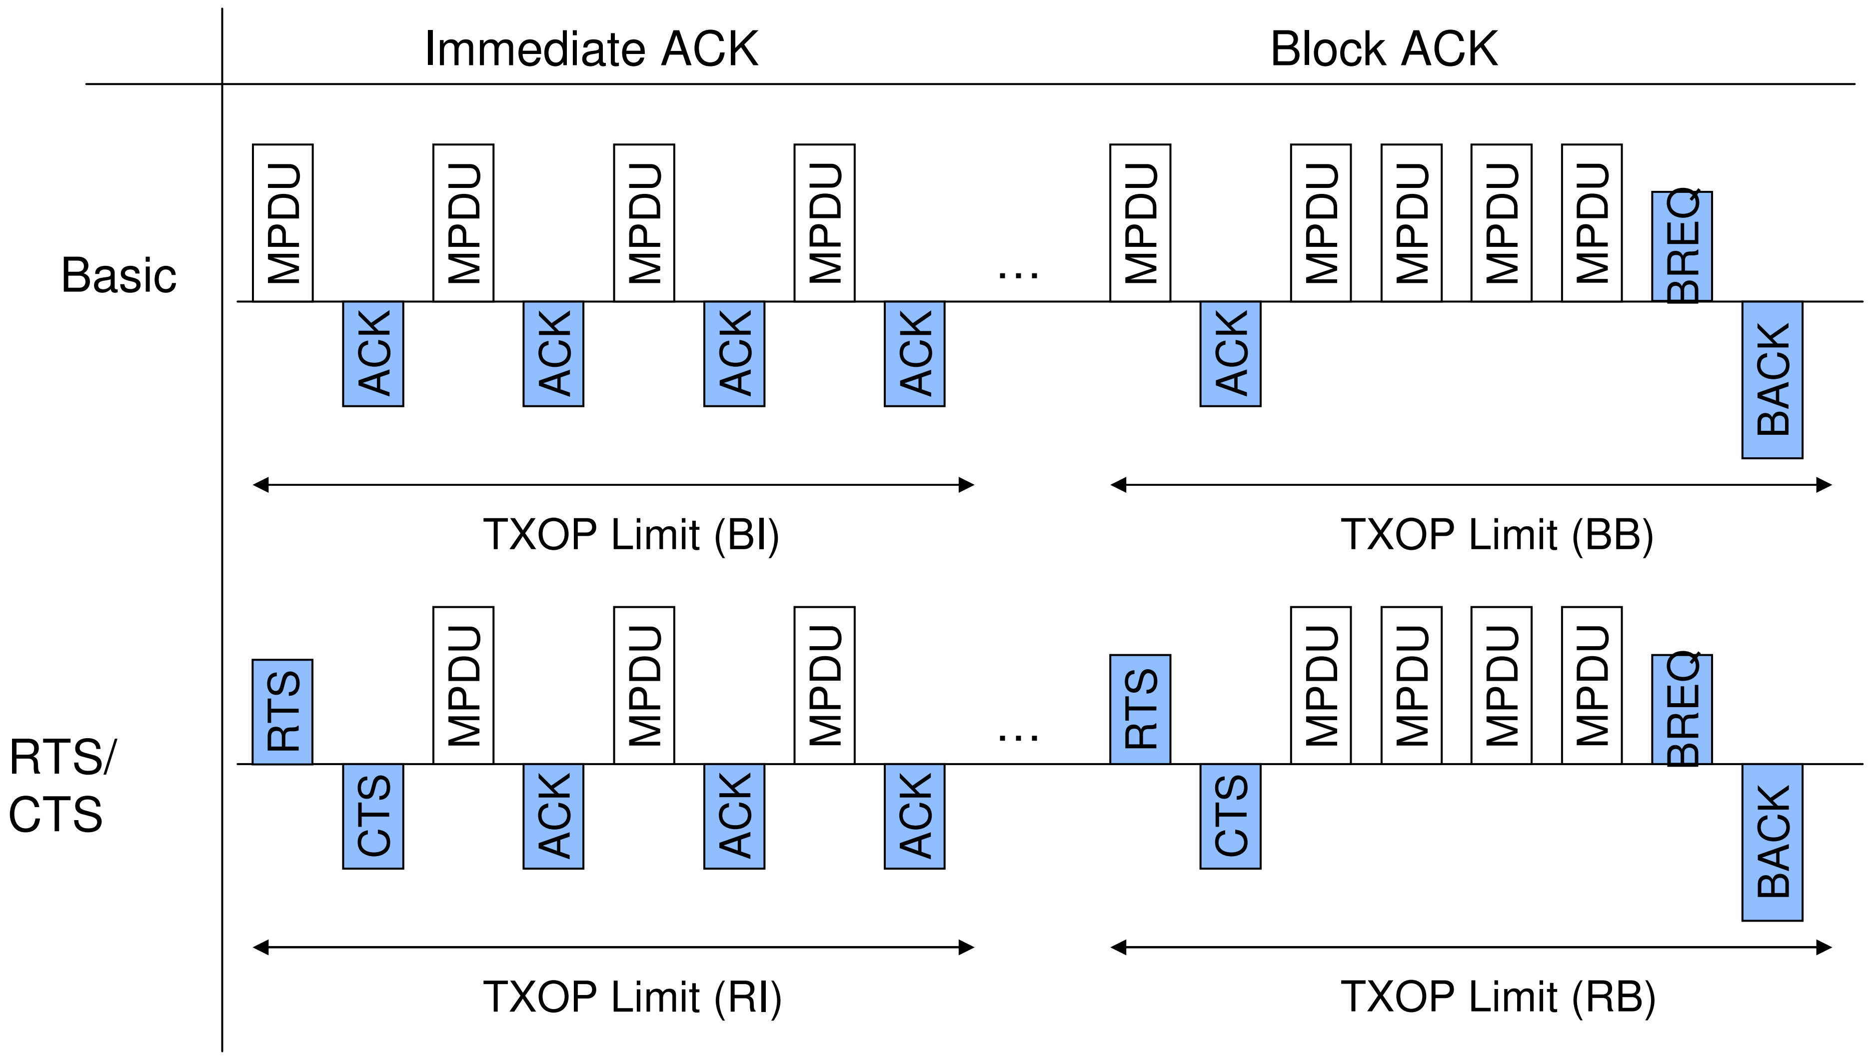

More on TXOP..

Basically, limit the channel holding times of the competing stations in presence of delay-sensitive traffic

However, TXOP implications are much deeper..

The channel access is managed with a completely different perspective
The access unit is not the MSDU (as in standard DCF), but a temporal interval -> temporary channel-service establishment *with higher efficiencies*

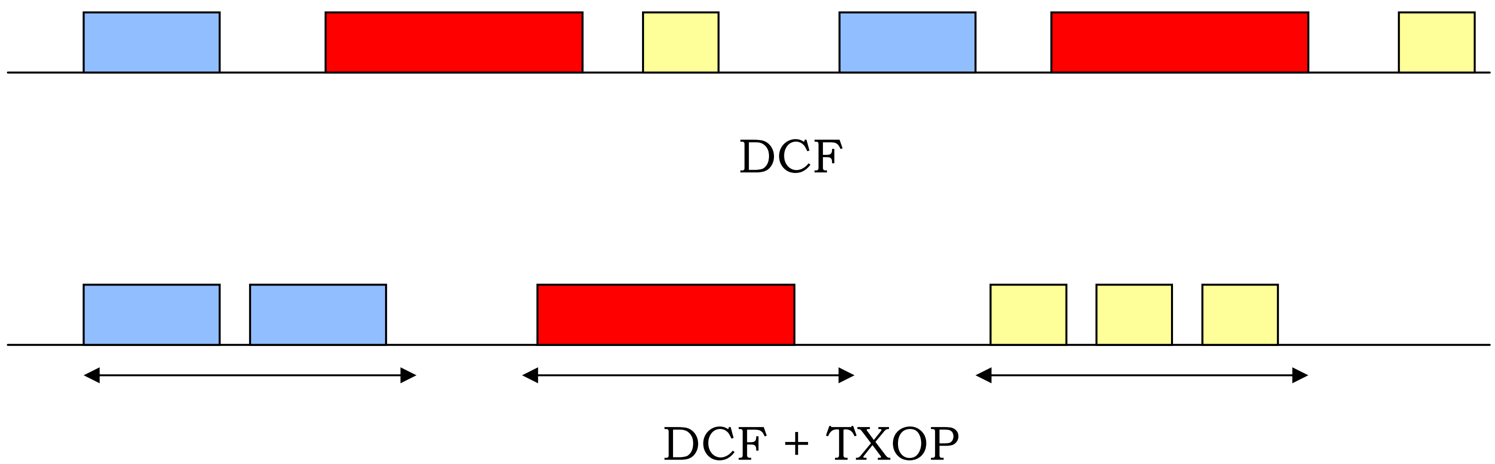

802.11e can natively provide temporal fairness via TXOP!