

CSE 461: Introduction to Computer Communications Networks Winter 2009

Module 3 Direct Link Networks – Part B

**John Zahorjan
zahorjan@cs.washington.edu
534 Allen Center**

Recap: We've Been Talking About Reliability

- Error correction
- ACKs + Retransmit
 - Are there other choices?
- Flow Control
 - Sliding window

This Module's Topics

Examples of Specific Protocols

1. Ethernet / IEEE 802.3
2. Wireless / IEEE 802.11

Relationship to the Protocol Stack


	Application	• Up to the application
	Presentation	• Encode/decode messages
	Session	• Manage connections
	Transport	• Reliability, congestion control
	Network	• Routing
	Link	• Framing, multiple access • Symbol coding, modulation
	Physical	

*Encoding
Framing*

- *Both protocols were designed to allow an evolving set of physical layer implementations (e.g., 802.11b, 802.11g, etc.).*
- *Link layer (can) consists of logical link control (LLC – multiplexing, reliability) and medium access control (MAC – framing, addressing, access control).*

Ethernet / 802.3

- Developed mid-1970's at Xerox PARC (along with the mouse, the bit-mapped display, and the personal workstation).
- Ethernet goals:
 - Cheap
 - Reliable
 - Passive network – a cable. No active elements required.
 - Taps are simple enough to be “fail stop.”
 - Distributed control – no central arbiter.
 - Statistical multiplexing.




(This shows the original, 3Mbps Ethernet. Modern versions typically look different at the physical level.)

We need...

- To define a policy for acquiring in the medium
 - Ethernet provides physical broadcast \Rightarrow want to avoid having more than one sender at a time
- To decide how much effort to put into reliability
 - Are all errors left to higher-layer protocols? Any?
- To define a frame format
 - What bits go where?
- To define an addressing scheme
 - How does sender name the intended receiver(s)?

Ethernet MAC



- The basic idea is to let a station send when it has something to send
 - We don't take turns; we don't make reservations
 - The Aloha network did just this
- The problem is that we'll have collisions, resulting in corrupted data
- We like to reduce the time wasted by collisions, without sacrificing distribution / simplicity

Reliability

- How much effort should be put into reliability at this level?
- Hey, it's a wire, with simple (reliable) attachments.
 - If there is no collision, the odds are very high the frame will be received correctly.
 - Don't embed into the MAC the overhead of acknowledgements
 - If there is a collision, the odds are not good that the frame will be received correctly.
 - There are no explicit ACKs, so assume the frame is lost in this case
- Upshot:
 - Send until there is no collision (or until you've tried enough times that you give up)

Reducing Collision Overheads – CSMA/CD

- Carrier Sense Multiple Access (CSMA)
 - Listen to medium before sending and “defer” if the medium is sensed busy
 - What collision scenarios does this eliminate?
 - Is it still possible to collide?
 - Is it likely?
- Collision Detection (CD)
 - A transmitting host also acts as a receiver
 - Detect a collision when the bit read from the ether is not the same as the one being transmitted
 - If a collision is detected, jam the ether so that all stations know there is a collision
 - What are the benefits of CD?

Implementing CSMA


- What happens when current transmission ends?
 - “1-persistent”: go ahead and send
 - “p-persistent”: a slotted time version of non-persistent
 - Each slot, assuming no transmission is on-going, start transmitting with probability p
 - Ideal: if N hosts want have data to transmit, p should be $1/N$
- By this definition, Ethernet is 1-persistent
 - It’s “optimistic” – let’s assume I’m the only station waiting to transmit ($N=1$)
- We’ll see that it also has a mechanism that can be thought of as attempting to dynamically estimate N (and to set $p=1/N$)

Reacting to Collision Detection

- What happens when a collision occurs?
 - Ethernet takes a collision as a sign it has underestimated N (the number of contending stations)
- Binary exponential backoff
 - A slot time is the maximum possible time between a host starting a transmission and all others hearing it
 - (Obviously, this is a function of the length of the Ethernet)
 - If k consecutive collisions have occurred, pick at random a number of slots between 0 and 2^k-1 and backoff (wait) that long before trying again
 - Binary exponential backoff has been proven stable in an idealized model
 - If all of N stations always have something to send, useful utilization of the Ethernet goes to $1/e$ as N gets large


Some Side Issues

- How does frame length affect efficiency?
- Why is there a minimum frame size?
 - How is it related to the maximum length of an Ethernet?
 - How is it related to the bandwidth?
- Modern Ethernets (100Mbps / 1Gbps / 10Gbps) look like this:



Why?

The Ethernet Frame / Addressing



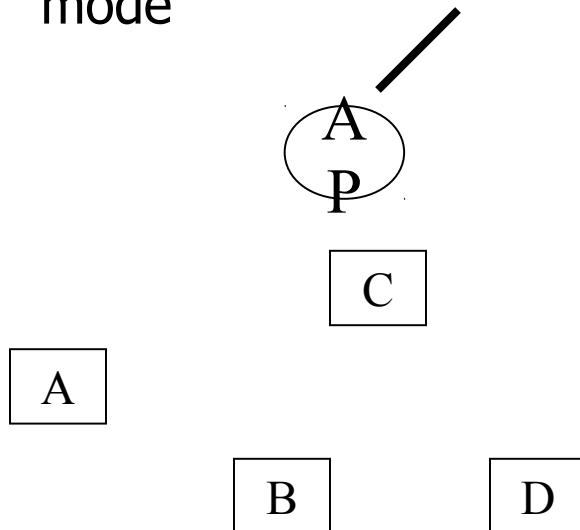
- Preamble lets the receiver synch
- Addresses are 6-bytes
- Type field allows demultiplexing
 - Overloaded to be a length field in some modern variants
- Minimum payload is 46-bytes; max is 1500
 - Pad is necessary if the actual data < 46 bytes
- You know what CRC is..

Ethernet (802.2) Addresses

- Each interface on an Ethernet has a unique address
 - Interface cards examine each frame as it goes by
 - If the destination address matches their own address, they save the frame and notify the host
 - (Interfaces can also be put into “promiscuous mode,” where they save all frames)
- Moreover, each interface in the world has a unique address
- Addresses are 48 bits ,written as sixteen hex digits
 - First 24 bits (4 million possibilities) identify a manufacturer (e.g., 3Com)
 - Last 24 bits are assigned by the manufacturer, so that all cards are unique
 - FF:FF:FF:FF:FF:FF is reserved as the broadcast address
- (Can you imagine other ways to assign addresses? Why is the one used attractive?)


Wireless / 802.11

- There is a lot of activity in the 802.11 world...
- We'll consider here
 - 802.11b (up to 11Mbps), 802.11a (up to 54Mbps), 802.11g (up to 54Mbps) [802.11n (up to 300Mbps)]
 - Distributed Coordination Function (DCF) / infrastructure mode



- All frames to/from a host go through the AP
- AP is connected to a larger network (e.g., the Internet) and acts as a relay

802.11 Wireless Networks



- Frequency division multiplexing is used statically
 - Each AP is on a channel (e.g., 802.11b has 13 channels)
- APs (typically) broadcast their service set ids (SSIDs)
- Clients select an AP and associate with it
 - Association has a medium term lifetime – many, many frames, typically
- Access to channels is through statistical multiplexing
- How should this work?

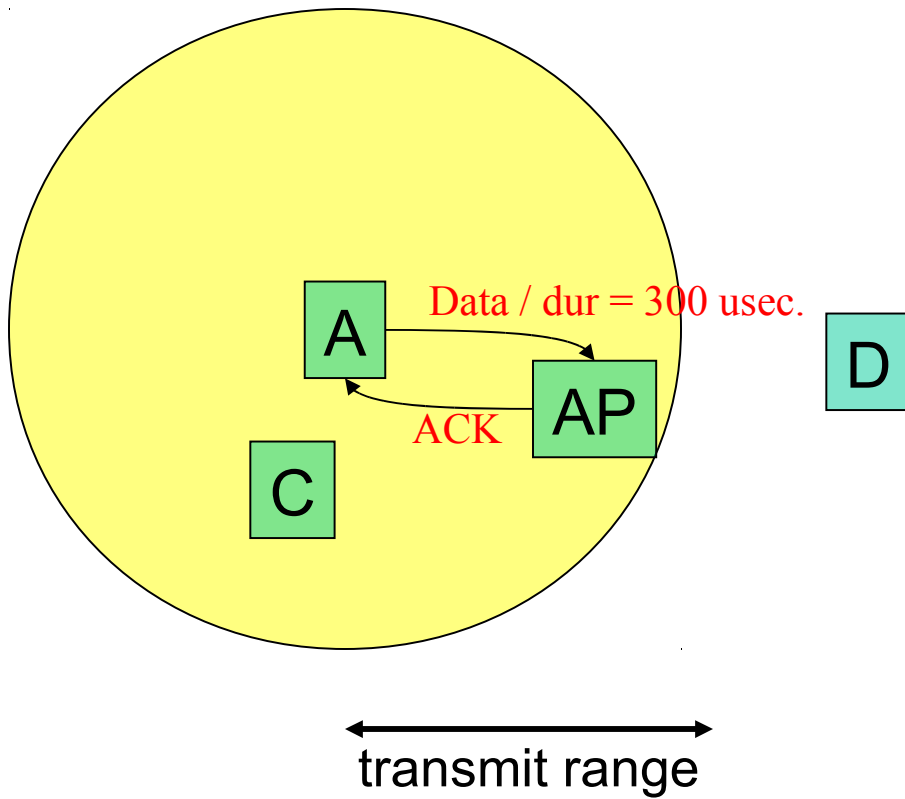
Characteristics of Wireless

- The ability of the radio to correctly decode a frame is determined by the signal-to-interference-and-noise-ratio (SINR):
 - $(\text{received signal strength}) / (\text{interference} + \text{noise}) > \beta$
- The received signal strength is the transmitted strength attenuated by the materials the signal passes through, and affected by multipath
- A useful but very inaccurate model is
 - $\text{received strength} = \text{sent strength} * d^{-\alpha}$
 - $\alpha=2$ for free space
 - $\alpha=3$ to 4 for in-building
- Interference is the energy of other on-going transmissions
- Noise is the energy generated by the receiving radio and other nearby sources (e.g., the computer's power supply)

Wireless Reliability

- Unlike Ethernet, frames can be lost even if only one station is transmitting
 - in fact, that's common
- 802.11 uses explicit receiver ACKs
- Time is "reserved" by each (data) frame for the ACK that should be coming back
 - (data) frames contain a duration field in the header
 - The duration is the time it will take to send the current frame, plus a short idle time, plus the time to send back the ACK
 - All stations hearing the current frame are required to remain silent until the duration time has elapsed

A Picture



Basic MAC Protocol

- Carrier-sense
 - Defer if you sense a sufficiently high energy level in the air
- No collision detection
 - Transmission emanating from radio overwhelms any incoming signal
- Explicit ACKs
 - If no ACK received in reserved time
 - Use a binary exponential backoff procedure to chose a random backoff time
 - Count down that time, pausing whenever you sense a transmission in the air
 - Re-transmit when your counter reaches 0

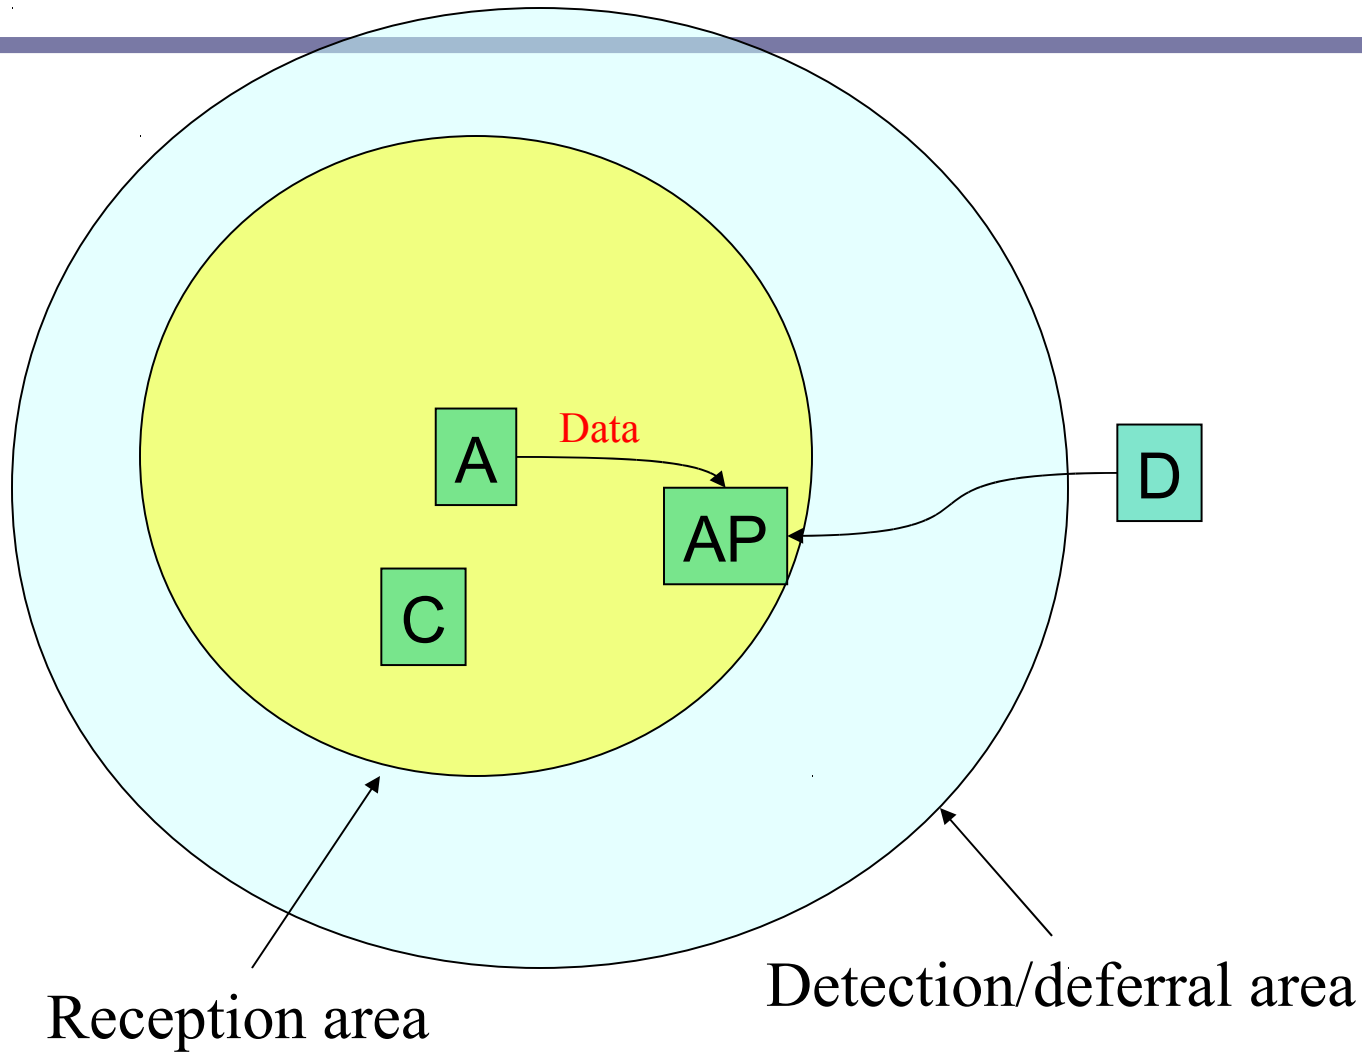
Modified ARQ

- To support this MAC level retry, the frame headers carry sequence numbers and a retry bit
 - Retry bit = 0 for first transmission of a frame, 1 for retries
 - Sequence number of each distinct frame must be distinct (until wraparound)
 - Allows receiver to detect (and throw away) duplicates
 - Same sequence number as last frame received from that source and retry bit = 1 means it's a duplicate
 - Otherwise, assume it's a new frame and deliver up to other protocol layers
- NO concept / detection of missing frames
 - Sequence numbers are used only to detect duplicates
 - "Missing" sequence numbers have no meaning
 - Successive sequence numbers to a particular destination may be any number not used too recently.

ACK Reliability

- Both the original frame and the ACK must be received or re-transmissions will take place
- As we've seen, time is reserved for the ACK, to help increase the odds it is received
- Additionally, ACKs are transmitted at the lowest rates
 - Multiple transmission rates are supported
 - E.g., 802.11b has 1, 2, 5.5, and 11Mbps
 - Slower rates have a lower SINR ratio for correct decoding

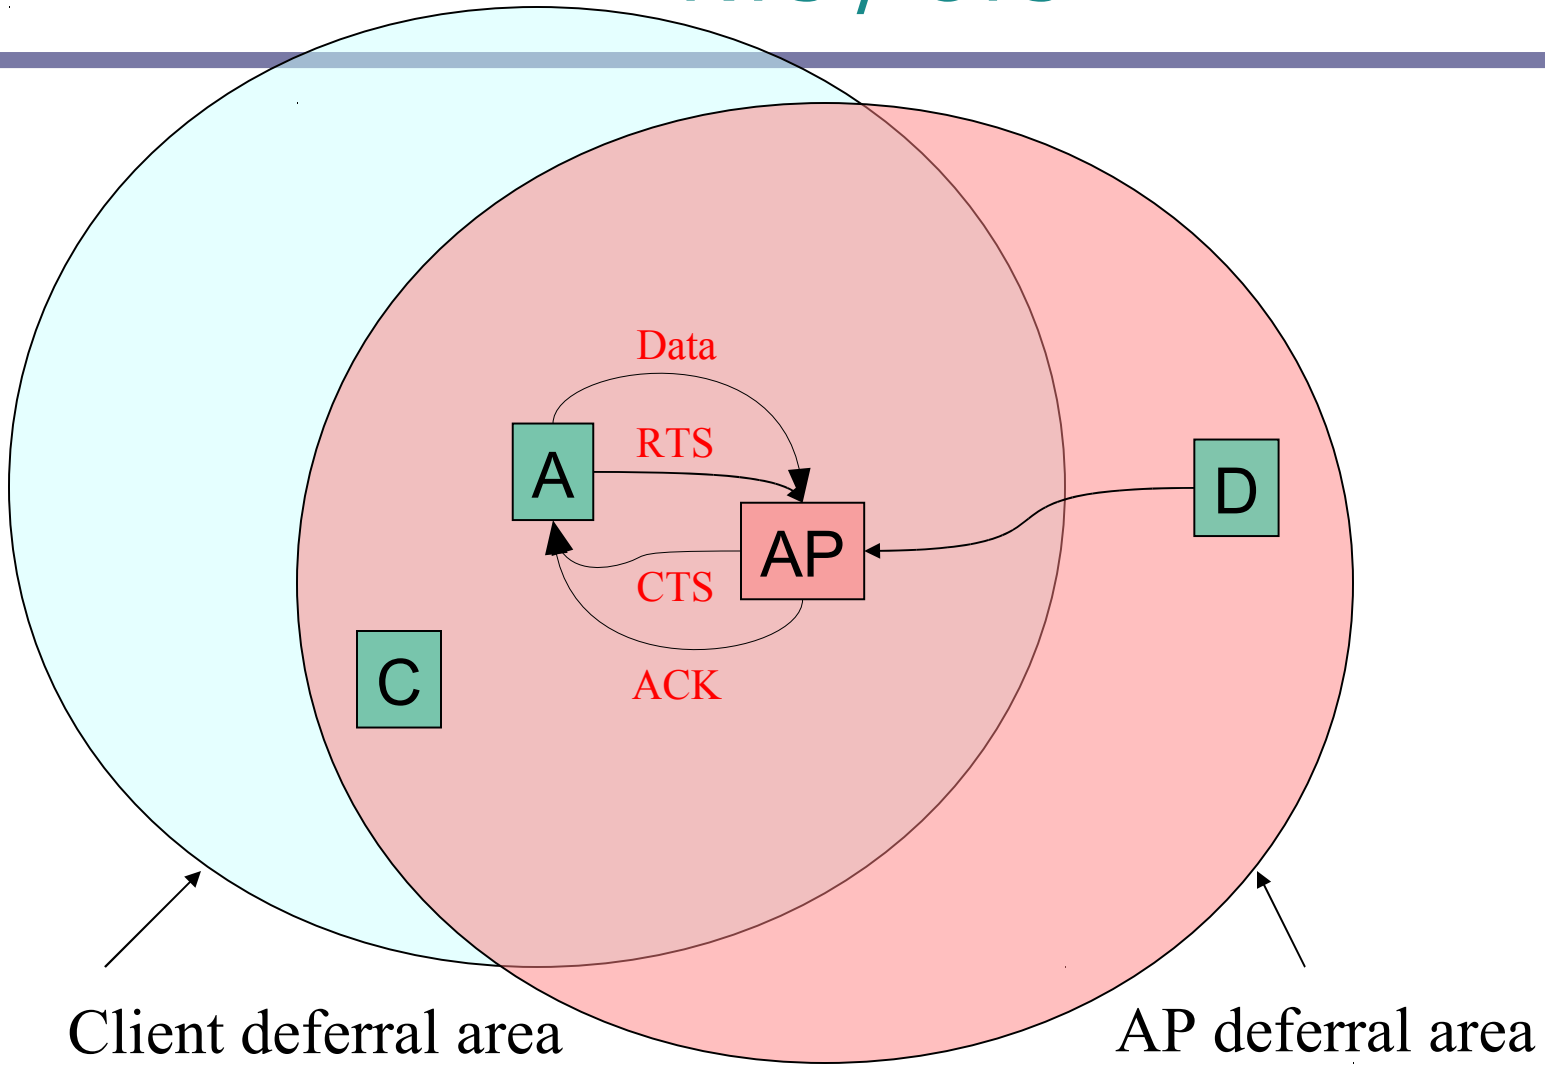
The Hidden Terminal Problem



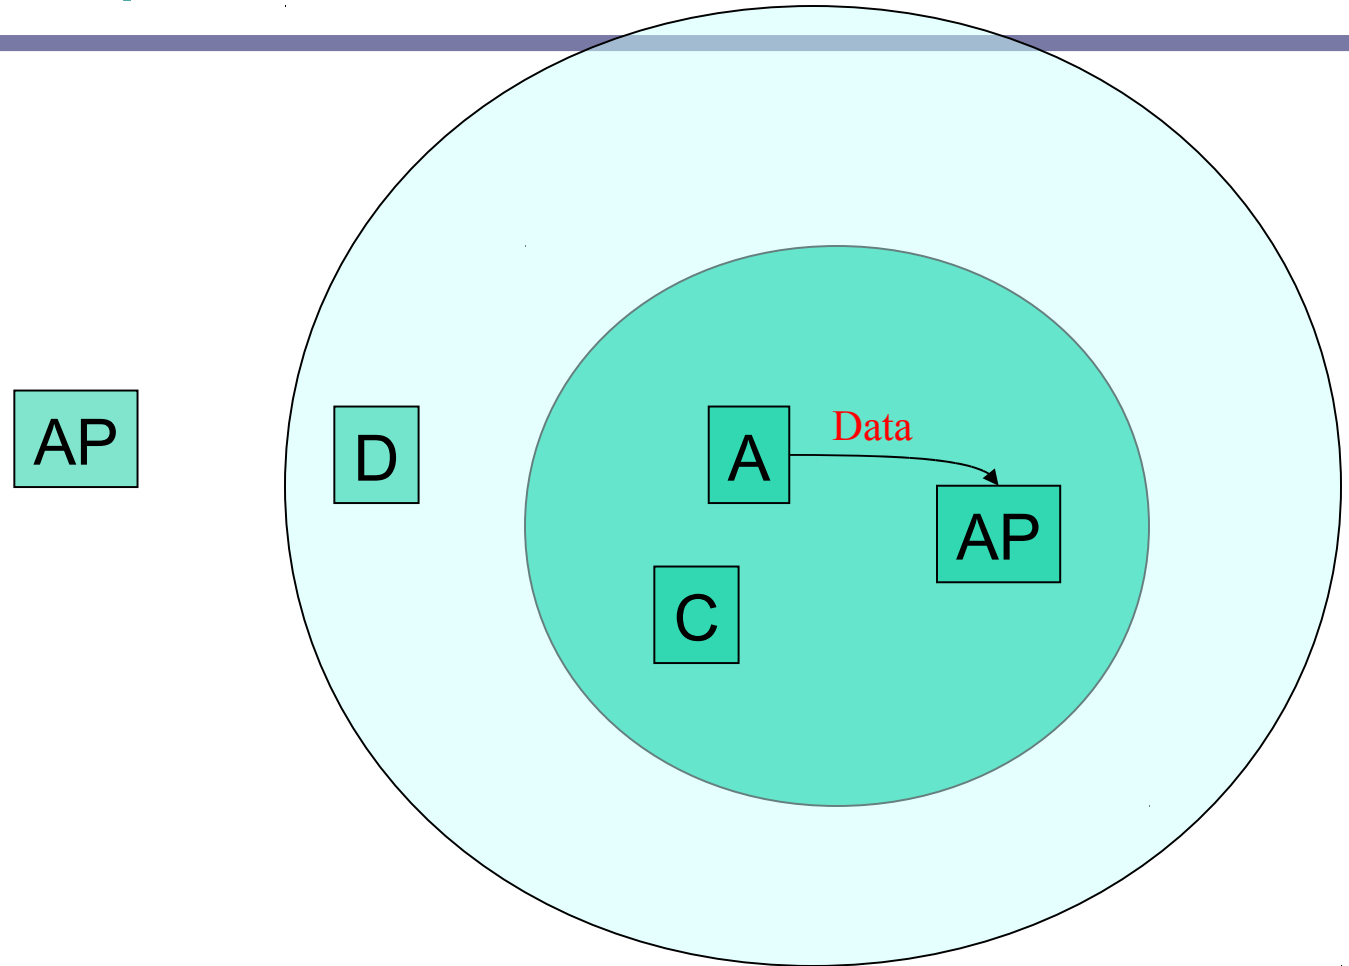
RTS / CTS

- A source may precede a data/ACK exchange with a request-to-send/clear-to-send (RTS/CTS) exchange
- The RTS carries a duration sufficient to cover the 4 frame exchange
 - With luck, it's heard by all other stations within range of the source
- The receiver responds with a CTS carrying the time required to cover the CTS / data / ACK
 - With luck, it's heard by all stations within range of the receiver
- If the CTS comes back, the source sends the data, in the normal way
- The specification does not dictate when to use RTS/CTS
 - It's actually much less used than the book implies
 - typically, there is a large, static frame size threshold, with RTS/CTS always used for frames larger than the threshold and never for those below

RTS / CTS



The Exposed Terminal Problem



Addressing

- 802.2 (48-bit) addresses are used
 - They're assigned just like with Ethernet – 24 bits name manufacturer, then 24 bits assigned by the manufacturer to that card
- Up to four addresses are contained in the header
 - Source: the address of where the frame originated
 - Transmitter: the address of the station actually transmitting
 - E.g., the AP might be forwarding a frame
 - Destination: address of the ultimate destination
 - Target: address of station that should take the frame off the air (e.g., the AP)

Frame Format

- Frames begin with a special bit pattern, sent at a low rate
- A zealous attempt has been made to keep frames as small as possible, leading to many frame types
- Here is a general idea of what they look like, though:




Figure 13—Frame Control field

Other Considerations

- There is an *ad hoc* mode, allowing stations to talk directly to each other (without the concept of an AP)
- The spec defines a *contention free* infrastructure (AP) mode in which the AP basically polls the clients for data
 - *This has perhaps never been implemented in any commodity hardware*
- There is support for *power management*
 - Clients may turn off their radios for a while
 - When they come back on, there are frame exchanges defined for them to ask for any frames the AP may be buffering for delivery

Summary

- Let's review the decisions made by Ethernet and 802.11:
 - Policy for acquiring in the medium
 - Ethernet: multiple access
 - 802.11: multiple access
 - How much effort to put into reliability
 - Ethernet: collision detect/ARQ; error detection/no ARQ
 - 802.11: error detection; ARQ (stop-and-wait)
 - To define an addressing scheme
 - Ethernet: source/dest 802.2 addresses
 - 802.11: 1-4 802.2 addresses
 - To define a frame format
 - Ethernet: max length (~1500 bytes); min length
 - 802.11: max length (~1500 bytes); many frame types/options (control and data frames)

A Quick Look At Alternatives

- Policy for acquiring in the medium
 - TDMA: wait your turn to send
 - Rings: a distinguished frame (“token”) is required before sending is allowed
 - Token is continuously circulated among all nodes
- How much effort to put into reliability
 - No error detection
 - Error correction (“forward error correction” (FEC))
 - Sliding-window?
- To define an addressing scheme
 - Hierarchical numeric addresses (rather than flat)?
 - Properties rather than addresses (“Canon iP6200 printer”)?
- To define a frame format
 - Larger frames?
 - Control frames / no control frames?