Les liens de communication et l'accès multiple

Bloc 2, INF 586

Walid Dabbous
INRIA Sophia Antipolis

Contenu du bloc

- La couche physique
  - notions de base sur la transmission
    - numérique
    - analogique
- L'accès multiple
  - les techniques de base
  - les protocoles
- Le contrôle de liaison
  - fera partie du 4ème bloc

Portée de ce bloc

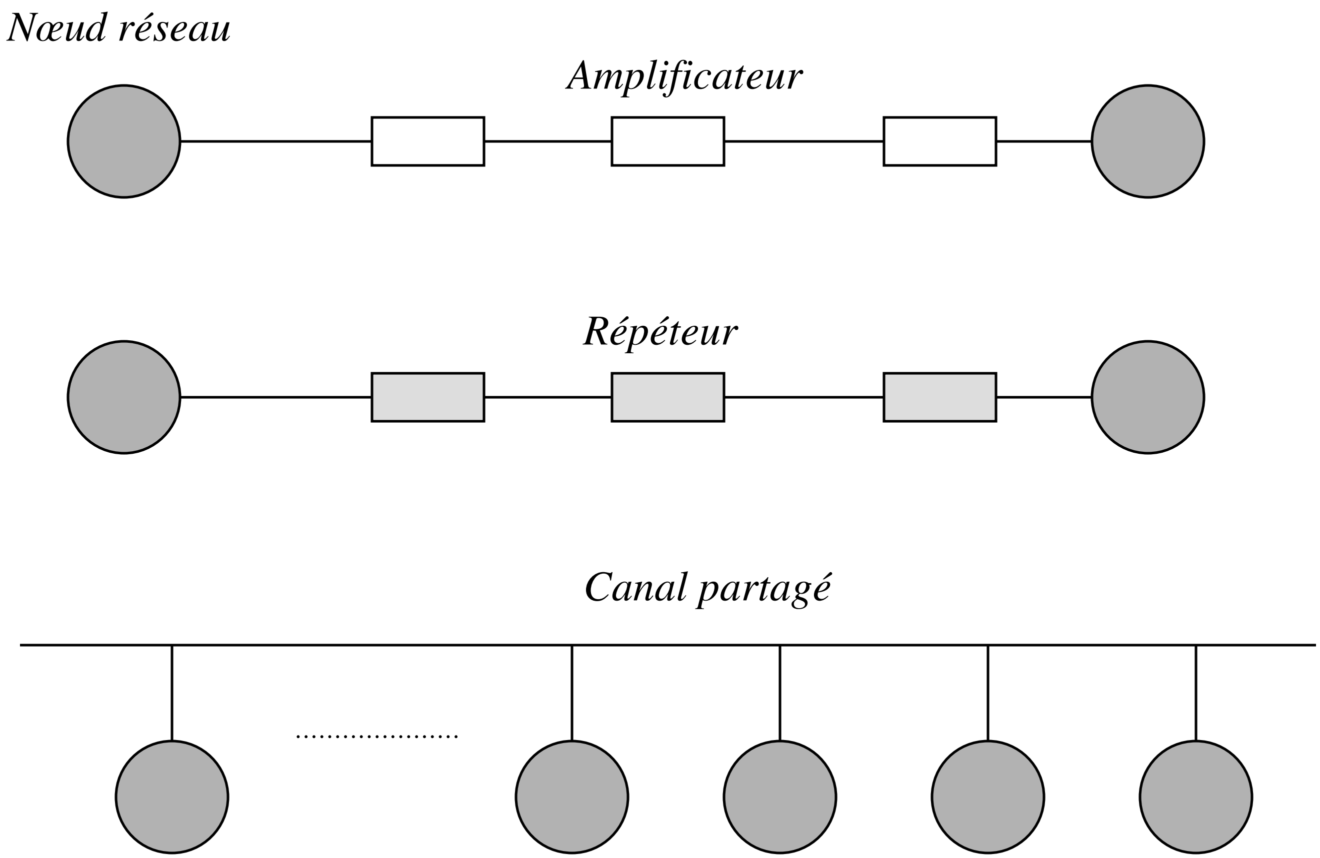

La couche physique

Transmission: some definitions

- Data: entities that convey information
- Analog: take continuous values in some interval
- Digital: take discrete values
- Analog data: voice, video, temperature, etc.
- Digital data: character strings, digitized audio or video, etc.
- Signal: Electric or electromagnetic representation of data
- Signaling (here): physical propagation of the signal along suitable medium
- Transmission: communication of data by the propagation and *processing* of signals

More definitions

- Analog signal: a continuously varying electromagnetic wave that may be propagated over a variety of (wired or wireless) media
- Digital signal: a discrete or discontinuous signal such as voltage pulses that may be propagated over a *wired* medium
- Wired media: guided transmission media e.g. twisted pair, coaxial cable, optical fiber (only transmits analog signal-encoded beam of light)
- Wireless (unguided) media: radio, terrestrial or satellite microwave, infrared
- Analog signaling: propagation of analog signal
- Digital signaling: propagation of digital signal
- Analog transmission: a means of transmitting analog signals without regard to their content (analog or digital data); amplifiers are used to boost the (analog) signal.

More definitions

- Digital transmission: the transmission of digital data, using either an analog or a digital signal, in which the digital data are recovered and repeated at intermediate points (repeaters) to reduce the effects of noise.
- coder (sampler+quantizer+digitize): transforms analog data/signal to a digital signal
- Modulator : transforms a digital bit stream into an analog signal suitable for transmission over analog media
- absolute bandwidth : width of the frequency spectrum ; $F_{max}-F_{min}$
- bandwidth: frequency band where « most » of the signal energy is contained (e.g. half-power bandwidth, 3dB below peak value)

Data rate and signal elements

- Signal element: the part of a signal that occupies the shortest interval of a signaling code
  - Digital signal element: a voltage pulse of constant amplitude
  - Analog signal element: a pulse of constant frequency, phase and amplitude
- Data element: a single binary one or zero (bit).
- Data rate (R in bps): rate at which data elements are transmitted
- Signaling or « modulation » rate (D in baud): rate at which signal elements are transmitted.
- Multilevel signal modulation techniques: reduces D because each signal element represents $b=Log_2$(number of levels) bits.

Relation between data rate and bandwidth (an example)

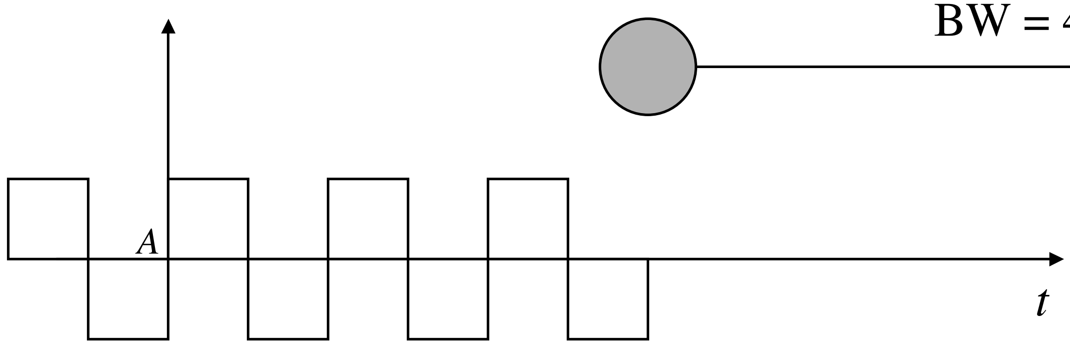

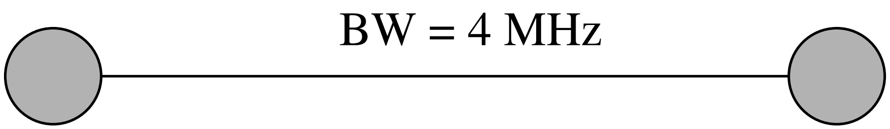

- Square wave s(t) with amplitudes A and -A, period T=1/f
- s(t) = A.4/π.Σ1/*k*.sin(2π*kf.t*) (*k* odd, *k*=1->inf)
- limit bandwidth to first three frequency components, T = 1 μs
  - Bw = 5MHz -1MHz=4MHz, R = 2 Mbps (1 bit every 0.5 μs)
- limit to first two frequency components, T=0.5 μs
  - Bw = 6MHz - 2MHz = 4MHz, R = 4Mbps (1 bit every 0.25 μs)
- Theortically unbounded, but economical and practical limitations -> distortion

Digital encoding techniques - comparison criteria

- Encoding scheme: mapping of data elements to signal elements
- Signal spectrum
  - lack of high frequency and dc components
  - concentrated spectral density
- Clocking
  - provide synchronisation with suitable encoding
- Error detection
  - built in error detection, faster than data link control
- signal interference and noise immunity
- Cost and complexity (increases with signalling rate)

Digital signal encoding formats

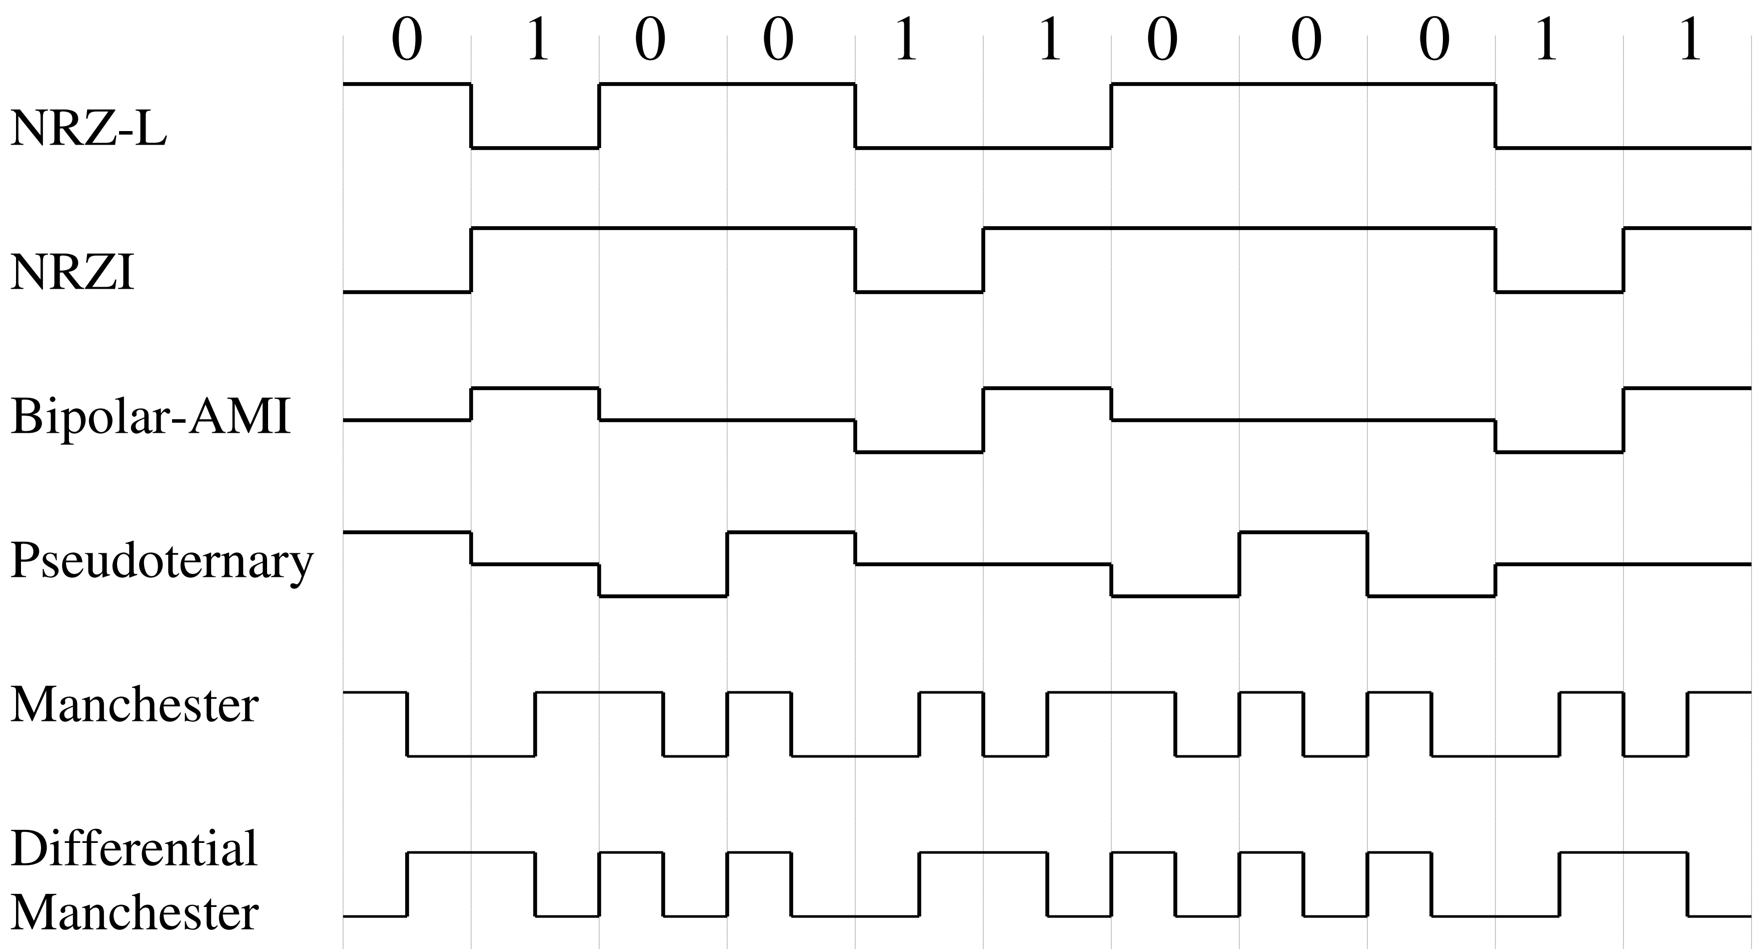

Encoding digital data to digital signals

- NonReturn to Zero-Level (NRZ-L): 0=+V, 1= -V
- NRZ Invert on 1s (NRZI): 0=no transition, 1=transition (differential encoding)
- Bipolar-AMI: 0=no line signal, 1= positive or negative level alternating for successive ones
- Pseudoternary: positive or negative level alternating for successive zeros, 1=no line signal
- Manchester: 0=transition from high to low in middle of bit interval, 1=transition from low to high
- Differential manchester (always a transition in middle of interval): 0: transition at beginning of interval, 1=no transition
- HDB3: based on Bipolar-AMI (four zeros replaced by code violation (signal pattern not allowed))

Encoding schemes characteristics

- NRZI uses differential encoding
  - to detect a transition is easier to compare to a threshold in presence of noise
  - immune to (accidental) polarity inversion
- NRZ codes are easy to enginneer and most energy is between DC and half the bit rate
  - but dc component and lack of synchronization
- Multilevel binary (Alternate Mark Inversion) (mark means 1, space means 0)
  - no dc component & provides better synchronization
    - however long string of 0s (AMI) or 1s (pseudoternary) still a pb
  - less efficient than NRZ (1 signal element carries 1 bit instead of $\log_2 3$ bits)
- Biphase
  - always a transition -> « self clocked » mechanism
  - no dc component, but bandwidth wider than multilevel binary

Encoding schemes characteristics

- Scrambling techniques (e.g. HDB3)
  - no dc component, no long sequence of zero-level line signals
  - no reduction in data rate (suitable for long distance transmission)
  - error detection capabilities
- Normalized signal transition rate of various schemes D =R/b

	Min. Mod Rate	101010…	Max. Mod Rate
NRZ-L	0 (all 0s or 1s)	1.0	1.0
NRZI	0 (all 0s)	0.5	1.0 (all 1s)
Bipolar-AMI	0 (all 0s)	1.0	1.0
Pseudoternary	0 (all 1s)	1.0	1.0
Manchester	1.0 (1010..)	1.0	2.0 (all 0s or 1s)
D-Manchester	1.0 (all 1s)	1.5	2.0 (all 0s)

Digital data to analog signals

- For transmission over the public telephone or microwave links
- Three encoding or modualtion techniques
- Amplitude shift keying (ASK):
  - $s(t) = A.\cos(2\pi f_c t)$ for binary 1, 0 for binary 0
- Frenquency shift keying (FSK):
  - $s(t) = A.\cos(2\pi f_1 t)$ for binary 1, $A.\cos(2\pi f_2 t)$ for binary 0
- Phase shift keying (PSK):
  - $s(t) = A.\cos(2\pi f_c t + \pi)$ for binary 1, $A.\cos(2\pi f_c t)$ for binary 0
- QPSK (signal element represents 2 bits)
  - $s(t) = A.\ \cos(b_1 b_2\, \pi).\cos(2\pi f_c t + 5\pi/4 + (b_2\text{-}b_1)\pi/2)$ for binary $b_1 b_2$

Digital transmission of analog data

- Digitize data then either:
  - transmit using NRZ-L
  - encode as a digital signal using code other than NRZ-L
  - convert digital data into an *analog* signal using modulation
    - allows digital transmission of voice on analog media e.g. microwave
- PCM
- Delta Modulation

Analog transmission

- Modulation for analog signals is useful
  - higher frequency needed for effective transmission (unguided transmission), baseband signals would require huge antennas
  - allows Frequency division multiplexing
- Input signal $m(t) = n_a . x(t)$
- Amplitude Modulation
  - $s(t) = [1 + m(t)] \cos(2\pi f_c t)$
- Angle modulation
  - $s(t) = A_c \cos[2\pi f_c t + \phi(t)]$
  - Phase modulation: $\phi(t) = n_p . m(t)$
  - Frequency modulation: $\phi'(t) = n_f . m(t)$

L’accès multiple

What is it all about?

- Consider an audioconference where
  - if one person speaks, all can hear
  - if more than one person speaks at the same time, both voices are garbled
- How should participants coordinate actions so that
  - the number of messages exchanged per second is *maximized*
  - time spent waiting for a chance to speak is *minimized*
- This is the *multiple access problem*

Some simple solutions

- A (simple) centralized solution: use a moderator
  - a speaker must wait for moderator to “poll” him or her, even if no one else wants to speak
  - what if the moderator’s connection breaks?
- Distributed solution
  - speak if no one else is speaking
  - but if two speakers are waiting for a third to finish, guaranteed collision
- Designing good schemes is surprisingly hard!

Outline

- Contexts for the problem
- Choices and constraints
- Performance metrics
- Base technologies
- Centralized schemes
- Distributed schemes

Contexts for the multiple access problem

- *Broadcast* transmission medium
  - message from any "station" is received by all receivers in its listening area
- Colliding messages are garbled
- Goal
  - maximize message throughput
  - minimize mean waiting time
- Shows up in five main contexts

Contexts

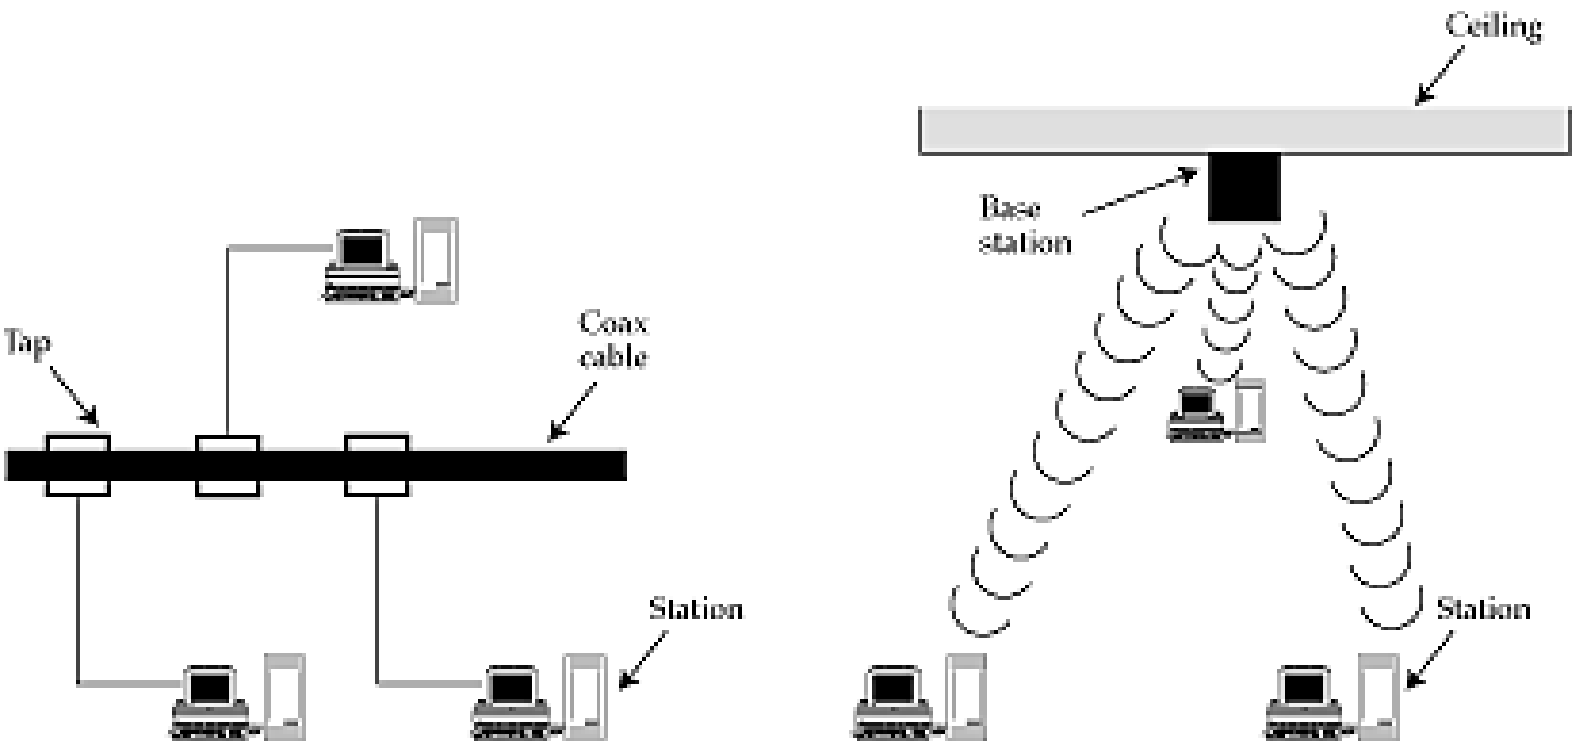

(a) Wired lan (b) Wireless lan

Contexts

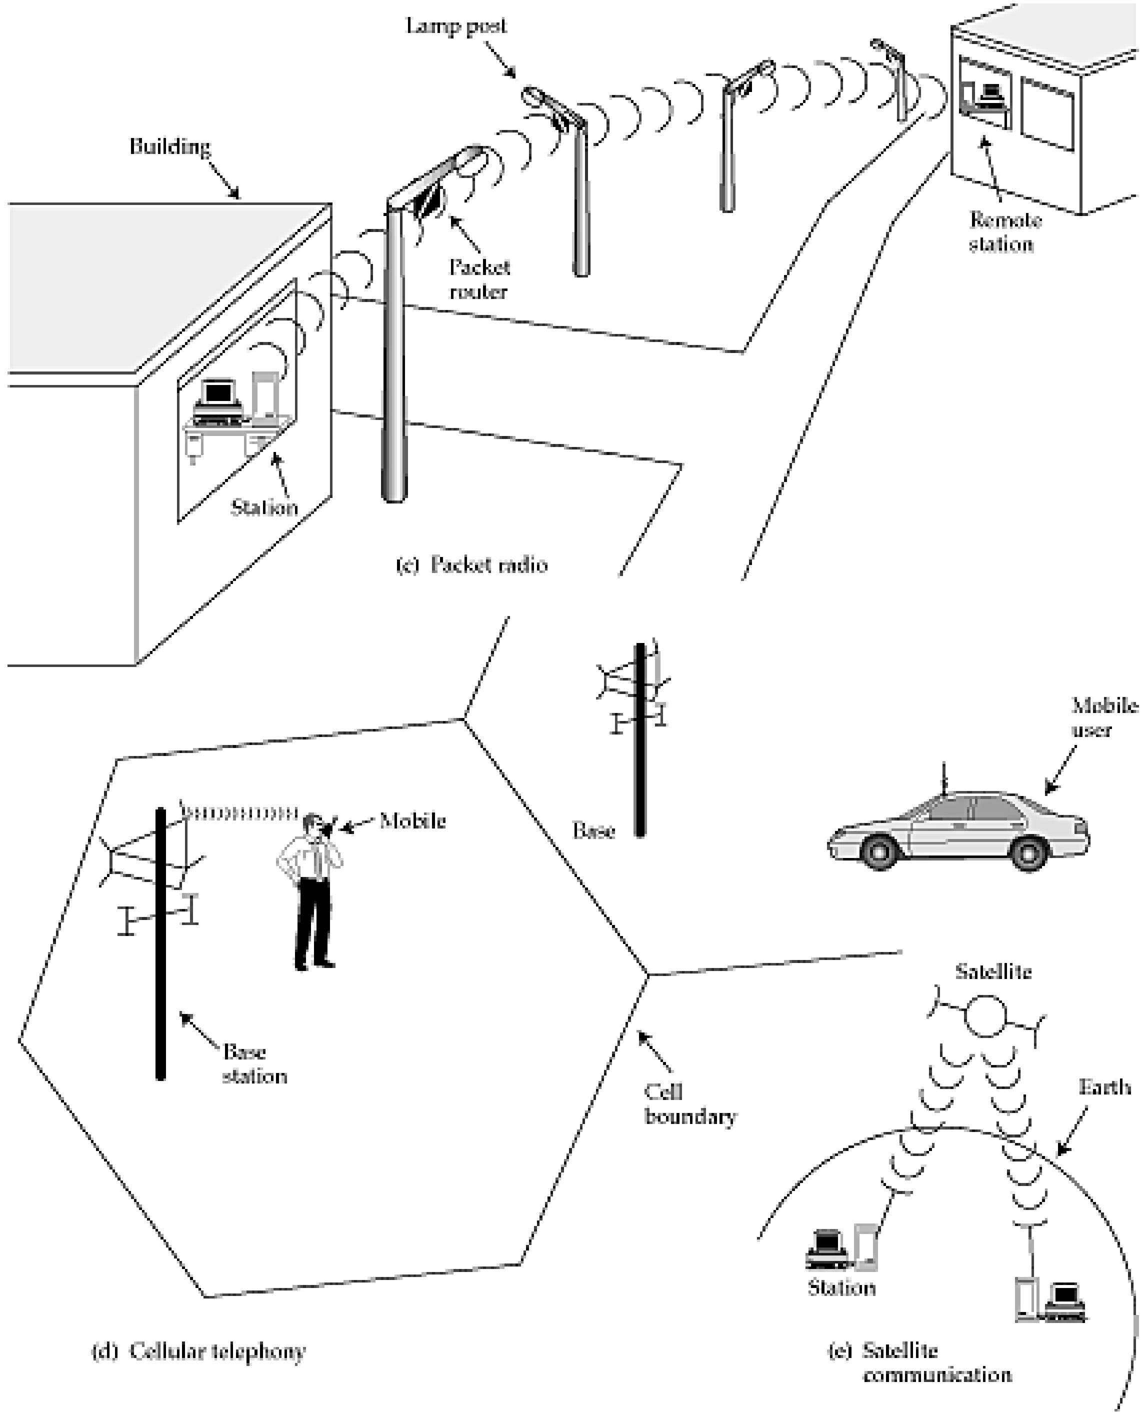
Lamp post
Building
Packet router
Remote station
Station
(c) Packet radio
Mobile
Base
Mobile user
Base station
Cell boundary
Satellite
Earth
Station
(d) Cellular telephony
(e) Satellite communication

Solving the problem

- First, choose a *base technology*
  - to isolate traffic from different stations
  - can be in time domain or frequency domain
- But: not enough time slots or frequencies to exclusively dedicate to each user
- Then, choose how to allocate a limited number of transmission resources to a larger set of contending users

Outline

- Contexts for the problem
- Choices and constraints
- Performance metrics
- Base technologies
- Centralized schemes
- Distributed schemes

Choices

- Centralized vs. distributed design
  - is there a moderator or not?
  - in a centralized solution one of the stations is a *master* (e.g. base station in cellular telephony) and the others are *slaves*
    - master->slave = downlink
    - slave->master = uplink
  - in a distributed solution, all stations are peers
- Circuit-mode vs. packet-mode
  - do stations send steady streams or bursts of packets?
  - with streams, doesn't make sense to contend for every packet
  - allocate resources to streams - analogy to long "speeches"
  - with packets, makes sense to contend for every packet to avoid wasting bandwidth - analogy to brainstorming session

Constraints

- Spectrum scarcity
  - radio spectrum is hard to come by
  - only a few frequencies available for long-distance (few miles) data communication
  - multiple access schemes must be careful not to waste bandwidth
- Radio link properties
  - radio links are error prone
    - fading (signal degradation because of hills, foliage, trucks, etc.)
    - multipath interference
  - hidden terminals
    - transmitter heard only by a subset of receivers
  - capture
    - station with higher power overpowers the other -> no “collision” is detected
    - lower powered station may never get a chance to be heard

The parameter ‘a’

- The number of packets sent by a source before the farthest station receives the first bit
- D: max propagation delay between any two stations
- T: time taken to transmit an average size packet
- a = D/T (around 10^{-2} for LANs, 1 for HSLANs, 10^{2} for satellite links)

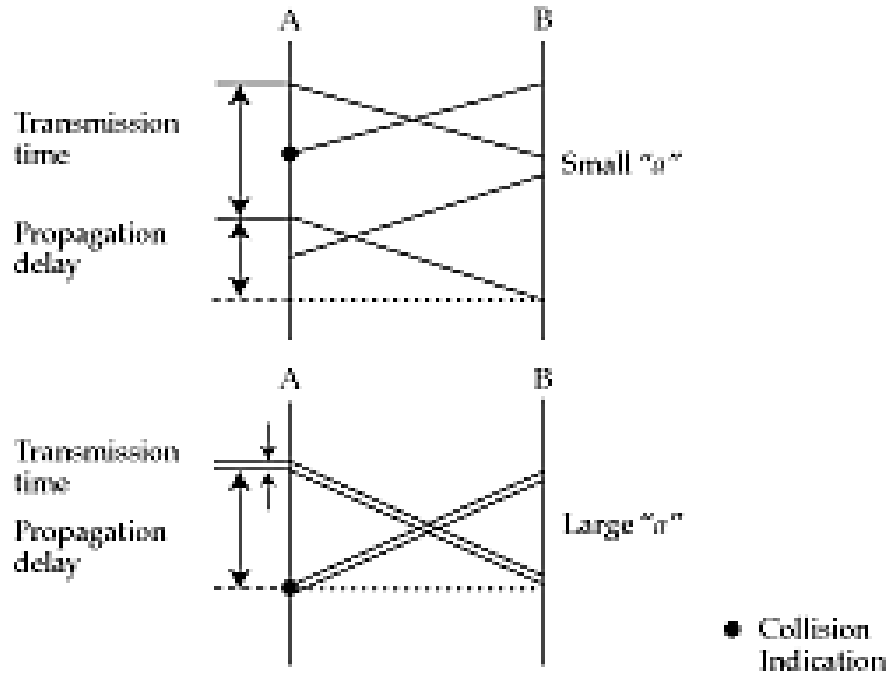

Outline

- Contexts for the problem
- Choices and constraints
- Performance metrics
- Base technologies
- Centralized schemes
- Distributed schemes

Performance metrics

- Normalized throughput
  - fraction of link capacity used to carry non-retransmitted packets
  - example
    - Ideally (with no collisions) a 1 Mbps link can carry 1000 “125 byte” packets/sec
    - with a particular scheme and workload, we might have 250 packets/sec
    - => goodput = 0.25
- Mean delay
  - amount of time a station has to wait before it successfully transmits a packet
    - depends on the MAC protocol, the load and the characteristics of the medium

Performance metrics

- Stability
  - with heavy load, is all the time spent on resolving contentions?
  - => unstable
  - with a stable algorithm, throughput does not decrease with offered load
  - if huge number of “uncontrolled” stations share a link, then instability is guaranteed
  - but if sources reduce load when overload is detected, can achieve stability
- Fairness
  - no single definition
  - ‘no-starvation’: source eventually gets a chance to send
  - stricter metric: each station gets an equal share of the transmission bandwidth

Limitation of analytical modeling

- Assumptions on source traffic
  - work of queuing theory is only indicative and not an absolute metric
- Other factors may have greater impact than what models predict
  - battery strength
  - weather conditions
  - presence or absence of foliage or plants
  - workload different from assumptions
- Will only focus on *qualitative* description

Outline

- Contexts for the problem
- Choices and constraints
- Performance metrics
- Base technologies
- Centralized schemes
- Distributed schemes

Base technologies

- Isolates data from different sources
- Three basic choices
  - Frequency division multiple access (FDMA)
  - Time division multiple access (TDMA)
  - Code division multiple access (CDMA)

FDMA

- Simplest, best suited for analog links
- Each station has its own frequency band, separated by guard bands
- Receivers tune to the right frequency
- Common for TV and radio broadcast (at most few hundred transmitters in a listening area) - Not in wireless telephony
- Number of frequencies is limited
  - reduce transmitter power; reuse frequencies in non-adjacent cells requiring “complex” handoff
  - trade off complexity for increased number of users
- simplistic example: voice channel = 30 KHz
  - 833 channels in 25 MHz band
  - with 7-cell pattern, partition into 119 channels each
  - but with N cells in a city, can get 119N calls => win if pattern repeated

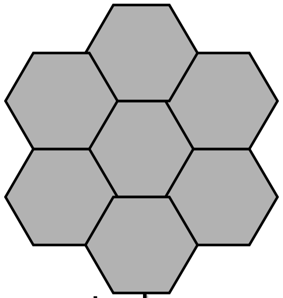

GSM frequency reuse

- Pattern: smallest cell set containing only once all radio channels (repeated on all zones to cover)
- Regular pattern size: 3,4,7,9, 12 & 21
- D_{ru} is re-use distance
- R is cell radius
- K is pattern size (7 in this example)
- $D_{ru}/R = sqrt(3.K)$

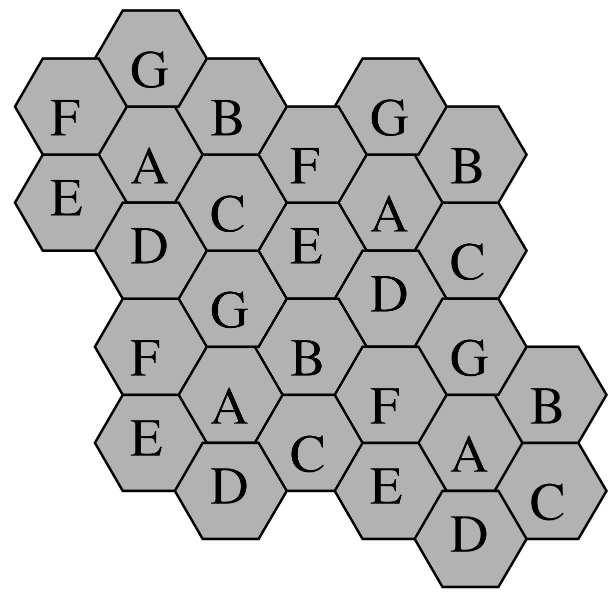

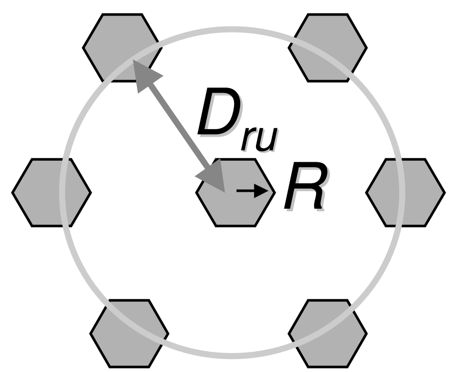

TDMA

- All stations transmit data on same frequency, but at different times
- Needs time synchronization
- supposes that stations resolve contention for access to a time *slot* and limit their transmission to a single slot
- roughly same number of users than FDMA
- Pros
  - users can be given different amounts of bandwidth (time slots)
  - mobiles can use idle times to determine best base station (and handoff to nearer base station to save battery)
  - can switch off power when not transmitting
- Cons
  - synchronization overhead (one of the stations emit sync signal)
  - greater problems with multipath interference on wireless links (because of wider frequency band -> smaller transmission duration -> need for adaptive equalization circuit in each receiver)

CDMA

- Users separated both by time and frequency
  - Aka spread spectrum techniques
- Colliding “frames” are not necessarily totally garbled
- Send at a different frequency at each time slot (*frequency hopping*)
- Or, convert a single bit to a code (*direct sequence*)
  - receiver can decipher bit by inverse process even in presence of narrowband noise
  - assumes that multiple signals add linearly

CDMA/DS example

- Each station has its « chip-sequence » or codeword (bipolar repsentation)
  - EA: (-1-1-1+1+1-1+1+1) (binary 00011011)
  - EB: (-1-1+1-1+1+1+1-1) (binary 00101110)
  - EC: (-1+1-1+1+1+1-1-1) (binary 01011100)
  - ED: (-1+1-1-1-1-1+1-1) (binary 01000010)
- Send codeword E for a 1 *bit* and its negation -E for a 0 *bit*

A	B	C	D
-	-	1	-
-	1	1	-
1	0	-	-
1	0	1	-
1	1	1	1
1	1	0	1

E1= (EC)	(-1 +1 -1 +1 +1 +1 -1 -1)
E2 = (EB+EC)	(-2 0 0 0 +2 +2 0 -2)
E3 = (EA-EB)	(0 0 -2 +2 0 -2 0 +2)
E4 = (EA-EB+EC)	(-1 +1 -3 +3 -1 -1 -1 -1 +1)
E5 = (EA+EB+EC+ED)	(-4 0 -2 0 +2 0 +2 -2)
E6 = (EA+EB-EC+ED)	(-2 -2 0 -2 0 -2 +4 0)

E1.EC	(1+1+1+1+1+1+1+1)/8 = 1
E2.EC	(2+0+0+0+2+2+0+2)/8 = 1
E3.EC	(0+0+2+2+0-2+0-2)/8 = 0
E4.EC	(1+1+3+3+1-1+1-1)/8 = 1
E5.EC	(4+0+2+0+2+0-2+2)/8 = 1
E6.EC	(2-2+0-2+0-2-4-0)/8 = -1

CDMA

- Pros
  - hard to spy
  - immune from narrow band noise
  - Unlike TDMA, no need for all stations to synchronize (only sender and receiver)
  - no hard limit on capacity of a cell, but noise increases and effective bit-rate per station decreases with the number of stations
  - all cells can use all frequencies, no need for frequency planning
- Cons
  - implementation complexity (receiver perfectly synchronized with senders)
  - need for complicated power control to avoid *capture*
    - hard because of multiple moving receivers
  - need for a large contiguous frequency band (for direct sequence) -> problems installing in the field

FDD and TDD

- Two ways of converting a wireless medium to a duplex channel
- In Frequency Division Duplex, uplink and downlink use different frequencies
- In Time Division Duplex, uplink and downlink use different time slots
- Can combine with FDMA/TDMA
- Examples
  - TDD/FDMA in second-generation cordless phones
  - FDD/TDMA/FDMA in digital cellular GSM phones

Outline

- Contexts for the problem
- Choices and constraints
- Performance metrics
- Base technologies
- Centralized schemes
- Distributed schemes

Centralized access schemes

- One station is master, and the other are slaves
  - slave can transmit only when master allows
- Natural fit in some situations
  - wireless LAN, where base station is the only station that can see everyone
  - cellular telephony, where base station is the only one
    - with a wired connection to the network and
    - capable of high transmit power

Centralized access schemes

- Pros
  - simple
  - master provides single point of coordination
- Cons
  - master is a single point of failure
    - need a re-election protocol => complicates the system!
    - master is involved in every single transfer => added delay
- Circuit mode schemes: cellular telephony
- Packet mode schemes: polling/probing and reservation

Circuit mode: the cellular telephony example

- When station wants to transmit, it sends a message to master using simple (ALOHA) packet mode multiple access protocol
- Master allocates transmission resources to slave
- Slave uses the resources until it is done
- No contention during data transfer
- Used primarily in cellular phone systems
  - EAMPS: analog FDD/FDMA
  - GSM : FDD/TDMA/FDMA
  - IS-95: CDMA

Polling and probing

- Packet-mode: station must contend for medium access for each packet
- Centralized controller mediates this contention
- Polling
  - master asks each station in turn if it wants to send (roll-call polling)
  - inefficient if (a) time to query a station is long, (b) overhead for polling messages is high, or (c) system has many terminals
- Probing
  - stations are numbered with consecutive logical addresses
  - assume station can listen both to its own address and to a set of multicast or “group” addresses
  - master does a binary search to locate next active station
    - skip chunks of address space with no active station
    - But repeated polls in sections with more than one active station
  - Efficient if few stations are active, doubles polls if all active

Reservation-based schemes

- When ‘a’ is large, collisions are expensive
  - polling overhead is too high
  - better use reservation than polling
  - mainly for satellite links
- Master coordinates access to link using reservations
- Some time slots devoted to reservation messages
  - can be smaller than data slots => *minislots*
- Stations contend for a minislot (PDAMA) (or own one FPODA)
- Master decides winners and grants them access to link
- Packet collisions are only for minislots, so overhead on contention is reduced

Outline

- Contexts for the problem
- Choices and constraints
- Performance metrics
- Base technologies
- Centralized schemes
- Distributed schemes

Distributed schemes

- Compared to a centralized scheme
  - more reliable
  - have lower message delays
  - often allow higher network utilization
  - but are more complicated
    - e.g. to synchronize stations to the same time base
- Almost all distributed schemes are packet mode
  - difficult to establish and maintain circuits without a central controller (the one-time coordination is amortized over many packets for circuit-mode in centralized schemes)

Decentralized polling

- Just like centralized polling, except there is no master
- But, all stations must share a time base
- Each station is assigned a slot that it uses
  - if nothing to send, slot is wasted
  - this is just TDMA :-)

Decentralized probing

- Also called *tree based multiple access*
- All stations in left subtree of root allowed to place packet on medium in first slot
- If a collision, root <- root ->left_son, and try again in next slot
- On success, everyone in root->right_son contend for access etc.
- Works well if ‘a’ is small
  - otherwise: either introduce idle time to wait for possible collision (inefficient) or roll back state if collision detected later (complex)

Carrier Sense Multiple Access (CSMA)

- Polling/probing may waste time if number of stations is large but number of simultaneously active stations is small
- A fundamental advance: check whether the medium is active before sending a packet (i.e *carrier sensing*)
- Unlike polling/probing a node with something to send doesn't have to wait for a master, or for its turn in a schedule
- If medium idle, then can send
  - just like a participant in a meeting
- If collision happens, detect and resolve
- Works when 'a' is small (0.1 or smaller)
- In slotted version, time *slot* is chosen to be the maximum *propagation* delay (considered small comparing to T)

Simplest CSMA scheme

- Send a packet as soon as medium becomes idle
- If, on sensing busy, wait for idle -> *persistent*
- If, on sensing busy, set a timer and try later -> *non-persistent*
- Problem with persistent: two stations waiting to speak will collide

How to solve the collision problem

- Two solutions
- *p-persistent*: when media becomes idle, transmit with probability *p:*
  - hard to choose *p* (< 1/Number of stations waiting)
    - balance message delay with higher utilization under heavy loads
  - if *p* small, then wasted time (if media idle)
  - if *p* large, more collisions
- *exponential backoff*
  - need to detect collisions: explicit ACK or *collision detect circuit* => CSMA/CD
  - on collision detection, choose retransmission timeout randomly from doubled range; on success reset timeout value
  - backoff range adapts to number of contending stations
  - no need to choose *p* (even 1-persistent CSMA with backoff is stable)

Summary of CSMA schemes

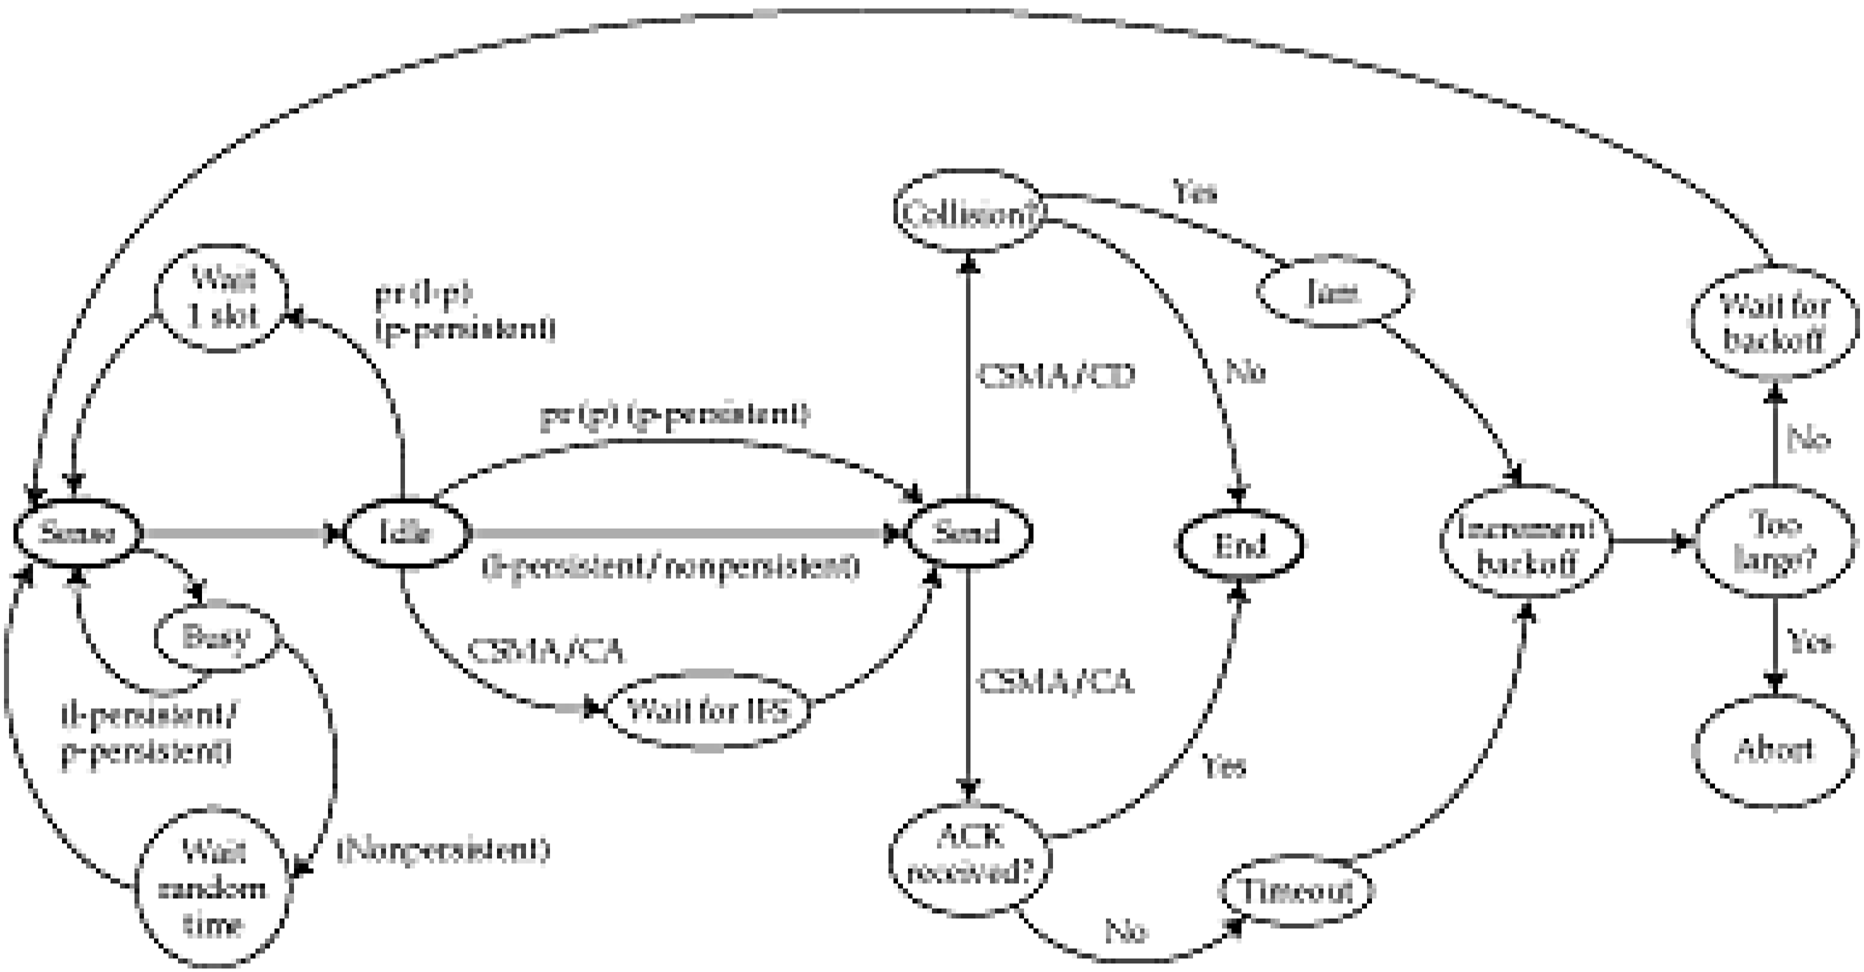

Ethernet

- The most widely used LAN
- Standard is called IEEE 802.3
- Uses 1-persistent CSMA/CD with exponential backoff
- Also, on collision, place a *jam* signal on wire, so that all stations are aware of collision and can increment backoff timeout range
- ‘a’ small =>time wasted in collision is around 50 microseconds
- Ethernet requires packet to be long enough that a collision is detected before packet transmission completes (a <= 1)
  - packet should be at least 64 bytes long for longest allowed segment
- Max packet size is 1500 bytes
  - prevents hogging by a single station

More on Ethernet

- First version ran at 3 Mbps
- Early versions ran at 10 Mbps, and uses ‘thick’ or ‘thin’ coax, or twisted pair
- Ethernet types are coded as <Speed><Baseband or broadband><physical medium>
  - Speed = 3, 10, 100 Mbps
  - Baseband = within building, broadband = on cable TV infrastructure
  - Physical medium:
    - “5” is thick coax, up to 500 meters
    - “2” is cheap 50 Ohm cable, up to 200 meters
    - “T” is unshielded twisted pair (also used for telephone wiring)

Enhancing Ethernet

- Ease of maintenance
  - use a hub as in 10BaseT
    - add/remove
- Increase performance
  - divide in multiple « contention domains »
    - use bridges
    - or (even) switches : Switched Ethernet
  - increase speed
    - 100BaseT

Some definitions

- Contention or collision domain: sum total of devices that compete with each other for access to the transmission media
- Hub: a centrally-located device in a star topology that propagates the signal transmitted by each node to ALL other ports. Nodes still constitute a single contention domain. Collision is detected by simultaneous activity on the Data Out (DO) and Receive Data (RD) circuits
- Bridge: device connecting « segments » with level 2 filtering capability. Splits LAN to N contention domains (N=number of ports). Packets are usually stored then forwarded.
- Switch: a « bridge » with N=number of nodes. If switch is full duplex capable, no collision will occur. Each pair gets 20Mbps

Recent developments

- Switched Ethernet
  - each station is connected to switch by a separate UTP wire
    - as in 10BaseT
  - however, line card of switch has a buffer to hold incoming packets
  - fast backplane switches packet from one line card to others
  - simultaneously arriving packets do not collide (until buffers overflow)
  - higher intrinsic capacity than 10BaseT (and more expensive)

Comparison

	Topology	Cable	Max distance	Nodes	Advantages	HalfDpX/FDX
10Base5	bus	Thick coax	500 x5seg	100/seg	backbones	HDX
10Base2	bus	Thin coax	185x5seg	30/seg	cheap	HDX
10BaseT	logical bus	2 UTPs	100 (Hub to Node)	few100s/CD	maintenance	HDX
Bridged	scalable	"bus" / port	No	per segment	multiple CDs	HDX in each CD
Switched	crossbar	UTPs	100 (btw swt.)	__	No contention	FDX

Fast Ethernet variants

- Fast Ethernet (IEEE 802.3u)
  - same as 10BaseT, except that line speed is 100 Mbps
  - spans only 205 m
  - big winner
  - most current cards support both 10 and 100 Mbps cards (10/100 cards) for about $80
- 100VG Anylan (IEEE 802.12)
  - station makes explicit service requests to master
  - master schedules requests, eliminating collisions
  - not a success in the market
- Gigabit Ethernet
  - aims to continue the trend
  - works over 4-pair UTP
- 10Gigabit Ethernet
  - No CSMA, only over optical fiber

Evaluating Ethernet

- Pros
  - easy to setup
  - requires no configuration
  - robust to noise
- Problems
  - at heavy loads, users see large delays because of backoff
  - non-deterministic service
  - doesn’t support priorities
  - big overhead on small packets
- But, very successful because
  - problems only at high load
    - loads rarely exceed 30%
  - can segment LANs to reduce load

CSMA/CA

- Used in wireless LANs
- Can’t detect collision because transmitter overwhelms colocated receiver
- So, need explicit acks
- But this makes collisions more expensive
  - => try to reduce number of collisions
- Standardized as IEEE 802.11

CSMA/CA algorithm

- First check if medium is busy
- If so, wait for medium to become idle
- if idle: wait for inter-frame spacing before contending for a slot (low IFS means higher priority)
- then, set a *contention timer* to an interval randomly chosen in the range [0, CW] (CW predefined contention window)
- On timeout, send packet and wait for ack
- If no ack, assume packet is lost
  - try again, after doubling CW
- If another station transmits while counting down, freeze CW and unfreeze when packet completes transmission
- station will get higher priority in next round of contention

Dealing with hidden terminals

- CSMA/CA works when every station can receive transmissions from every other station
- Not always true
- Hidden terminal
  - some stations in an area cannot hear transmissions from others, though base can hear both (C cannot sense A is sending to B)
- Exposed terminal
  - some (but not all) stations can hear transmissions from stations not in the local area (B should be able to send to D, while A sending to C)

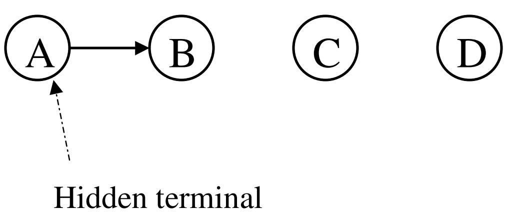

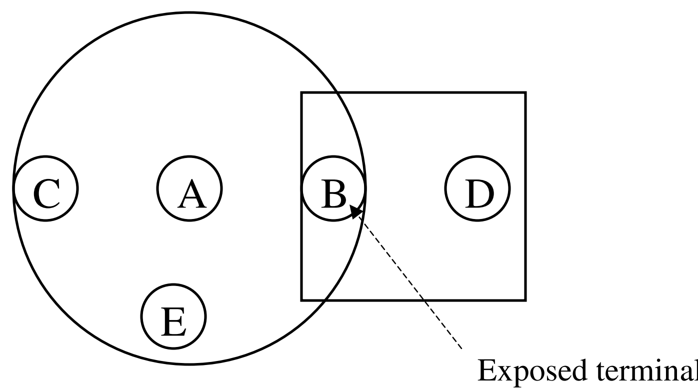

Dealing with hidden and exposed terminals

- In both cases, CSMA/CA doesn't work
  - with hidden terminal, collision because carrier not detected
  - with exposed terminal, idle station because carrier incorrectly detected
  - what matters is collision "at the receiver"
- Two solutions
- Busy Tone Multiple Access (BTMA)
  - assumes symmetric wireless links
  - uses a separate "busy-tone" channel
  - when station is receiving a message, it places a tone on this channel
  - everyone who might want to talk to a station knows that it is busy
    - even if they cannot hear transmission that this station hears
  - this avoids both problems of hidden and exposed terminals
    - transmitters ignore their carrier-sense circuit and sends only if busy-tone channel is idle

Multiple Access Collision Avoidance

- BTMA requires us to split frequency band
  - more complex receivers (need two tuners)
- Separate bands may have different propagation characteristics
  - scheme fails!
- Instead, use a single frequency band, but use explicit messages to tell others that receiver is busy
- In MACA, before sending data, send a Request to Sent (RTS) to intended receiver
- Station, if idle, sends Clear to Send (CTS)
- Sender then sends data
- If station overhears RTS, it waits for other transmission to end
- Solves both problems

Token passing

- In distributed polling, every station has to wait for its turn
- Time wasted because idle stations are still given a slot
- What if we can quickly skip past idle stations?
- This is the key idea of token ring
- Special packet called ‘token’ gives station the right to transmit data
  - analogy with “right to speak or microphone”
- When done, it passes token to ‘next’ station
  - => stations form a logical ring
- No station will starve
- In addition, stations no longer need precise time synchronization

Logical rings

- Can be on a non-ring physical topology

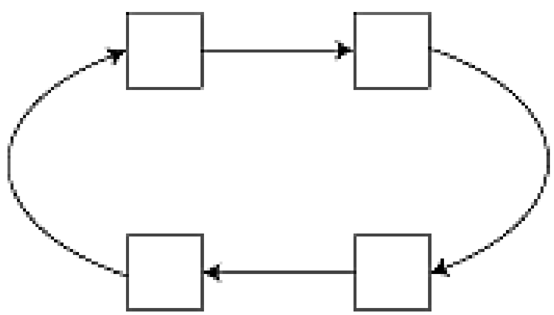

(a) Single ring

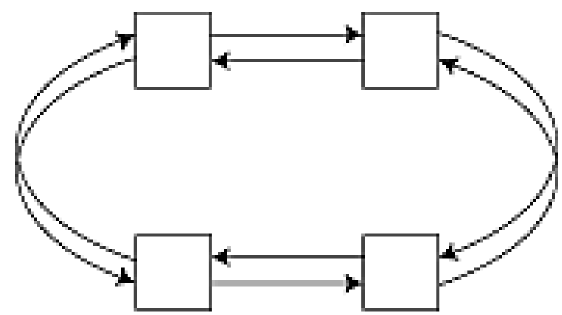

(b) Dual ring

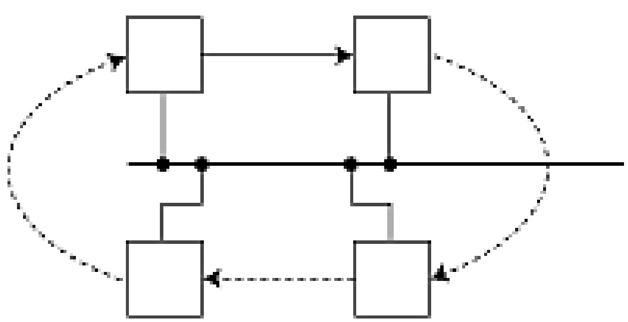

(c) Token bus

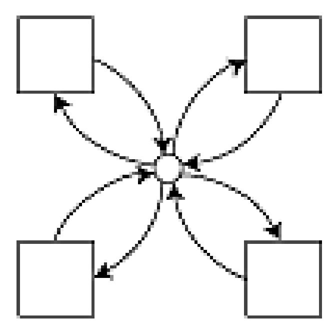

(d) Hub or
star-ring

Ring operation

- During normal operation, copy packets from input buffer to output
- If packet is a token, check if packets ready to send
- If not, forward token
- If so, delete token, and send packets
- Receiver copies packet and sets 'ack' flag
- Sender removes packet from the ring
- When done, reinserts token
- If ring idle and no token for a long time, regenerate token

Single and double rings

- With a single ring, a single failure of a link or station breaks the network => fragile
- With a double ring, on a failure, go into *wrap mode*
- Used in FDDI

Hub or star-ring

- Simplifies wiring
- Active hub is predecessor and successor to every station
  - can monitor ring for station and link failures
- Passive hub only serves as wiring concentrator
  - but provides a single test point
- Because of these benefits, hubs are practically the only form of wiring used in real networks
  - even for Ethernet

Evaluating token ring

- Pros
  - medium access protocol is simple and explicit
  - no need for carrier sensing, time synchronization or complex protocols to resolve contention
  - guarantees zero collisions
  - can give some stations priority over others
- Cons
  - token is a single point of failure
    - lost or corrupted token trashes network
    - need to carefully protect and, if necessary, regenerate token
  - all stations must cooperate
    - network must detect and cut off unresponsive stations
  - stations must actively monitor network
    - to detect token loss and duplication
    - usually elect one station as *monitor*

Fiber Distributed Data Interface

- FDDI is the most popular token-ring base LAN
- Dual counter-rotating rings, each at 100 Mbps
- Uses both copper (CDDI) and fiber links
- Supports both non-realtime and realtime traffic
  - token is guaranteed to rotate once every Target Token Rotation Time (TTRT)
  - station is guaranteed a *synchronous allocation* within every TTRT
- Supports both *single attached* and *dual attached* stations
  - single attached (cheaper) stations are connected to only one of the rings

ALOHA and its variants

- ALOHA is one of the earliest multiple access schemes
- Just send it!
- Wait for an ack
- If no ack, try again after a random waiting time
  - no backoff

Evaluating ALOHA

- Pros
  - useful when ‘a’ is large, so carrier sensing doesn’t help
    - satellite links
  - simple
    - no carrier sensing, no token, no timebase synchronization
  - independent of ‘a’
- Cons
  - under some mathematical assumptions, goodput is at most .18
    - much higher goodput is achievable (e.g. a single user Trx)
  - at high loads, collisions are very frequent
  - sudden burst of traffic can lead to instability
    - unless backoff is exponential

Slotted ALOHA

- A simple way to double ALOHA's capacity
- Make sure transmissions start on a slot boundary
  - synchronize to a broadcast pulse
- Halves *window of vulnerability*
- Used in cellular phone uplink (to request a frequency)
  - stations need to synchronize for TDMA anyway -> little additional overhead

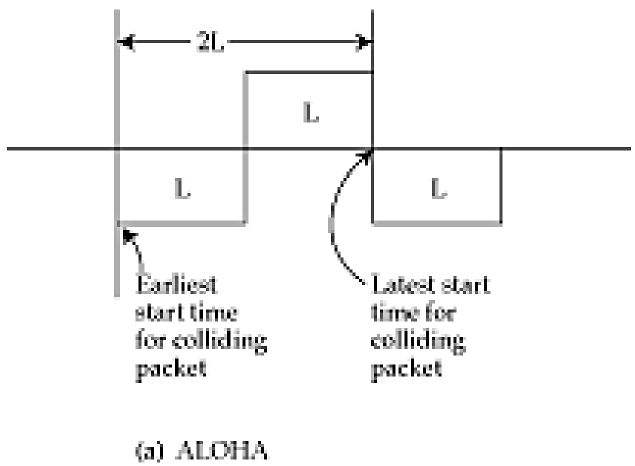

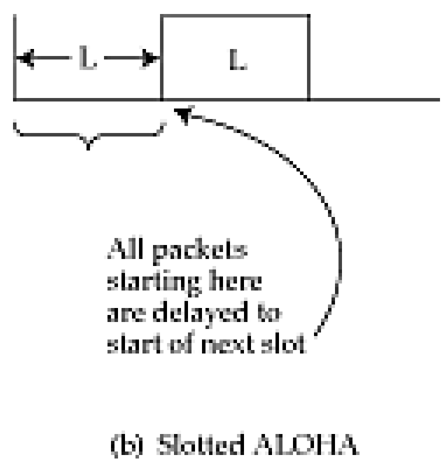

ALOHA schemes summarized

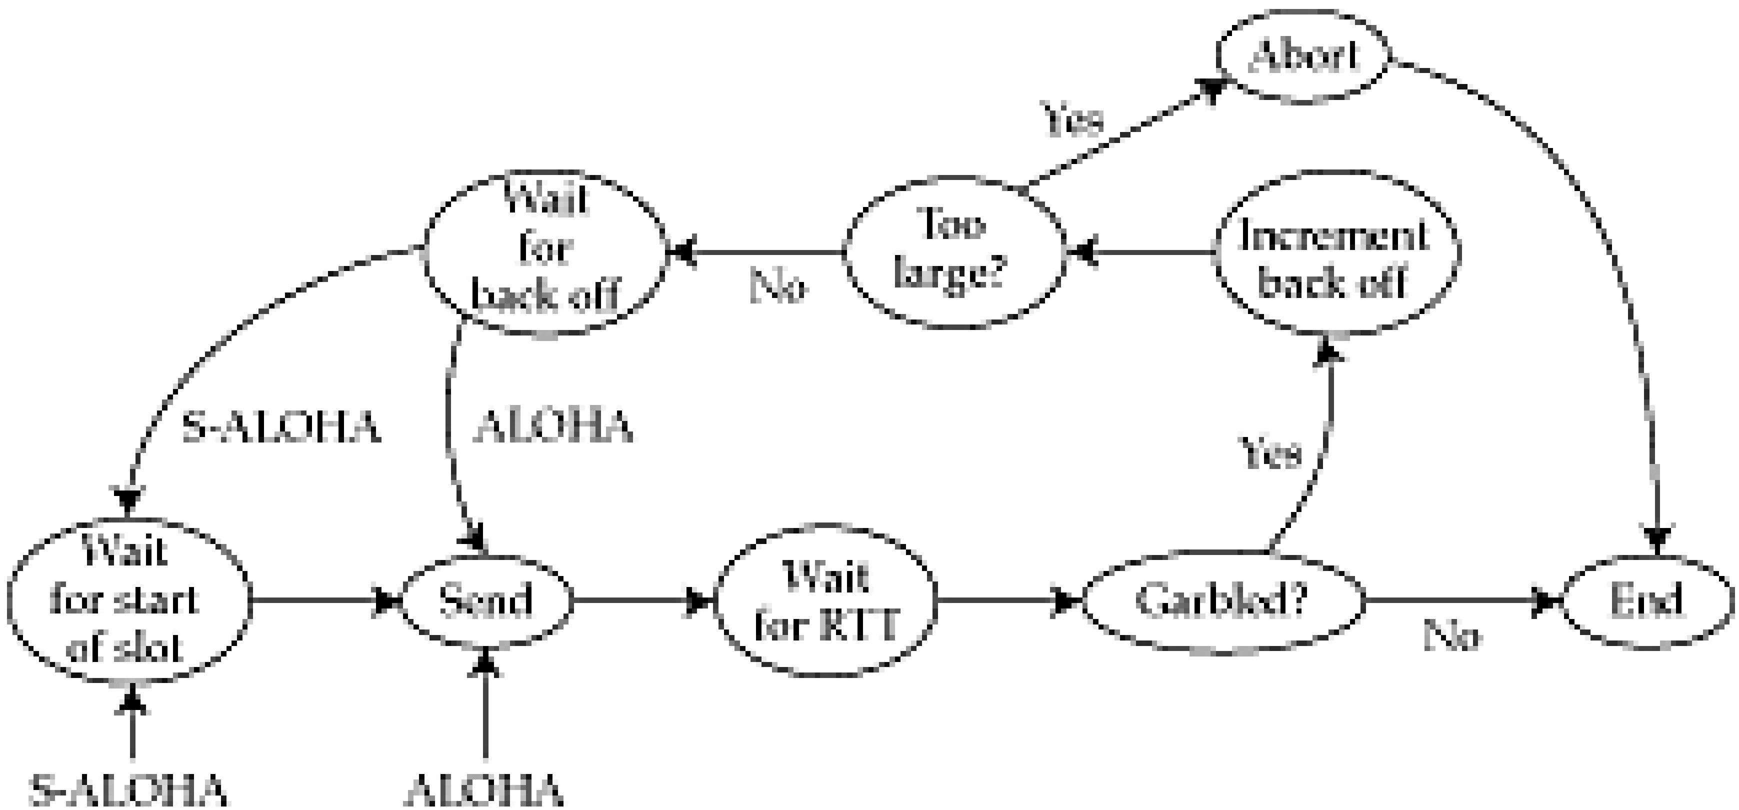

Reservation ALOHA

- Similar to FPODA but no master, slot time = time to transmit a (constant length) packet
- Contend for implicit reservation using slotted ALOHA
- Stations independently examine reservation requests and come to consistent conclusions
  - all stations have same priority
- Simplest version
  - divide time into frames = fixed length set of slots (spans prop delay)
  - station that successfully transmit in a slot using S-ALOHA (this is known implicitly in satellite Aloha and with explicit ACK with cellular Aloha) makes implicit reservation and can keep slot as long as it wants
  - station that loses keeps track of idle slots and contends for them in next frame

Evaluating R-ALOHA

- Pros
  - supports both circuit and packet mode transfer
  - works with large ‘a’
  - simple
- Cons
  - arriving packet has to wait for entire frame before it has a chance to send (frame length is at least ‘a’)
  - cannot preempt hogs
  - variants of R-ALOHA avoid these problems

Classification of multiple access schemes

Base technologies	FDMA				
	TDMA				
	CDMA				
	FDD				
	CDD				
Access schemes	Centralized	Circuit mode	EAMPS		Cellular telephony
			GSM		
			IS-95		
		Packet mode	Polling and probing		Wired LAN
			Reservation based	FPODA	satellite
				PDAMA	
	Distributed	Packet mode	Polling and probing		Wired LAN
			CSMA	CSMA/CD	
				CSMA/CA	Wireless LAN
			BTMA	MACA	
				MACAW	
			Token ring	FDDI	Wired LAN
			ALOHA	P-ALOHA	satellite
				S-ALOHA	
				R-ALOHA	