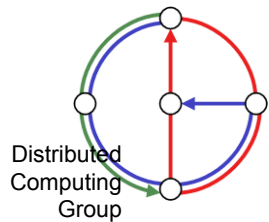


Chapter 3

MEDIA ACCESS CONTROL



Mobile Computing
Summer 2003

Overview



- Motivation
- SDMA, FDMA, TDMA
- Aloha
- Adaptive Aloha
- Backoff protocols
- Reservation schemes
- Polling



Motivation

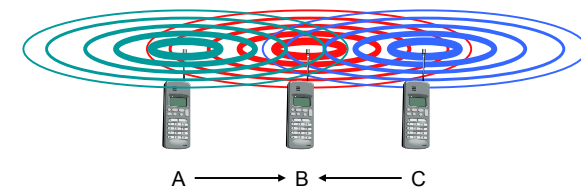


- Can we apply media access methods from fixed networks?
- Example CSMA/CD
 - **C**arrier **S**ense **M**ultiple **A**ccess with **C**ollision **D**etection
 - send as soon as the medium is free, listen into the medium if a collision occurs (original method in IEEE 802.3)
- Problems in wireless networks
 - signal strength decreases at least proportional to the square of the distance
 - senders apply CS and CD, but the collisions happen at receivers

Motivation – Hidden terminal problem

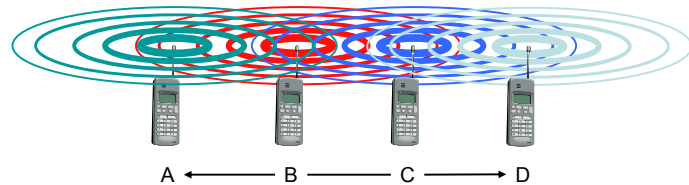


- A sends to B, C cannot receive A
- C wants to send to B, C senses a “free” medium (CS fails)
- collision at B, A cannot receive the collision (CD fails)
- A is “hidden” for C



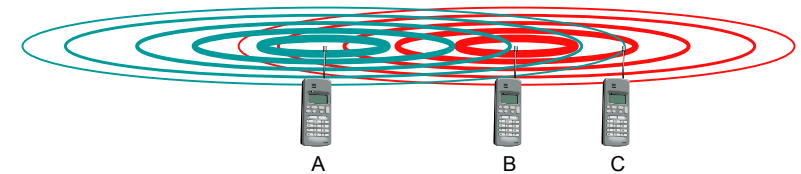
Motivation – Exposed terminal problem

- B sends to A, C wants to send to D
- C has to wait, CS signals a medium in use
- since A is outside the radio range of C waiting is not necessary
- C is “exposed” to B



Motivation - near and far terminals

- Terminals A and B send, C receives
 - the signal of terminal B hides A's signal
 - C cannot receive A



- This is also a severe problem for CDMA networks
- precise power control

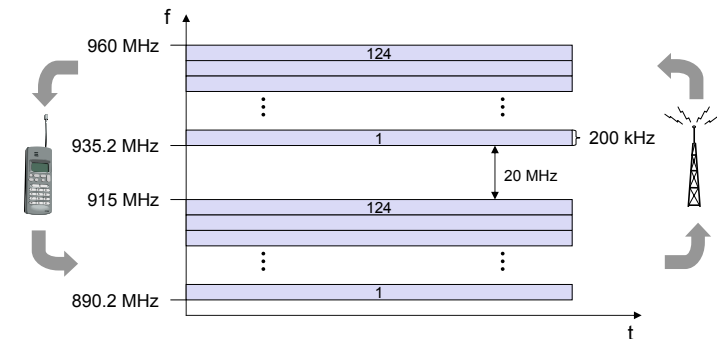


Access methods SDMA/FDMA/TDMA

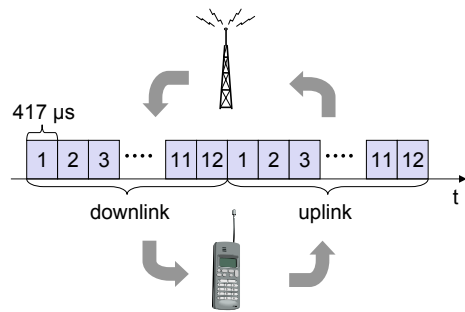
- SDMA (Space Division Multiple Access)
 - segment space into sectors, use directed antennas
 - Use cells to reuse frequencies
- FDMA (Frequency Division Multiple Access)
 - assign a certain frequency to a transmission channel
 - permanent (radio broadcast), slow hopping (GSM), fast hopping (FHSS, Frequency Hopping Spread Spectrum)
- TDMA (Time Division Multiple Access)
 - assign a fixed sending frequency for a certain amount of time



FDD/FDMA - general scheme, example GSM

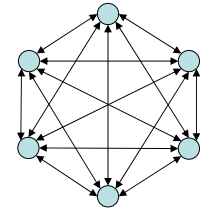


TDD/TDMA - general scheme, example DECT



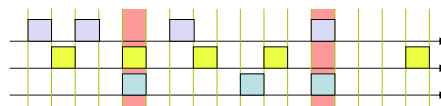
TDMA – Motivation

- We have a system with n stations $(0, 1, 2, \dots, n-1)$ and one shared channel
- The channel is a perfect broadcast channel, that is, if any single station transmits alone, the transmission can be received by every other station. There is no hidden or exposed terminal problem. If two or more transmit at the same time, the transmission is garbled.
- Round robin algorithm: station k sends after station $k-1 \pmod n$
- If a station does not need to transmit data, then it sends “ ϵ ”
- There is a maximum message size m that can be transmitted
- How efficient is round robin? What if a station breaks or leaves?
- All deterministic TDMA protocols have these (or worse) problems



TDMA – Slotted Aloha

- We assume that the stations are perfectly synchronous
- In each time slot each station transmits with probability p .



$$P_1 = \Pr[\text{Station 1 succeeds}] = p(1-p)^{n-1}$$

$$P = \Pr[\text{any Station succeeds}] = nP_1$$

$$\text{maximize } P : \frac{dP}{dp} = n(1-p)^{n-2}(1-pn) = 0 \Rightarrow pn = 1$$


$$\text{then, } P = \left(1 - \frac{1}{n}\right)^{n-1} \geq \frac{1}{e}$$

- In slotted aloha, a station can transmit successfully with probability at least $1/e$. How quickly can an application send packets to the radio transmission unit? This question is studied in queuing theory.



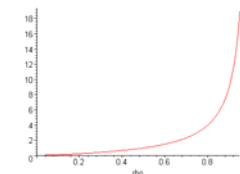
Queuing Theory – the basic basics in a nutshell

- Simplest M/M/1 queuing model (M=Markov):
- Poisson arrival rate λ , exponential service time with mean $1/\mu$



- In our time slot model, this means that the probability that a new packet is received by the buffer is λ ; the probability that sending succeeds is μ , for any time slot. To keep the queue bounded we need $\rho = \lambda/\mu < 1$.

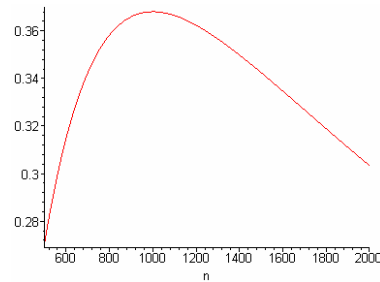
- In the equilibrium, the expected number of packets in the system is $N = \rho/(1-\rho)$, the average time in the system is $T = N/\lambda$.



Slotted Aloha vs. Round Robin

- Slotted aloha uses not every slot of the channel; the round robin protocol is better.
- + What happens in round robin when a new station joins? What about more than one new station? Slotted aloha is more flexible.

- Example: If the actual number of stations is twice as high as expected, there is still a successful transmission with probability 30%. If it is only half, 27% of the slots are used successfully.



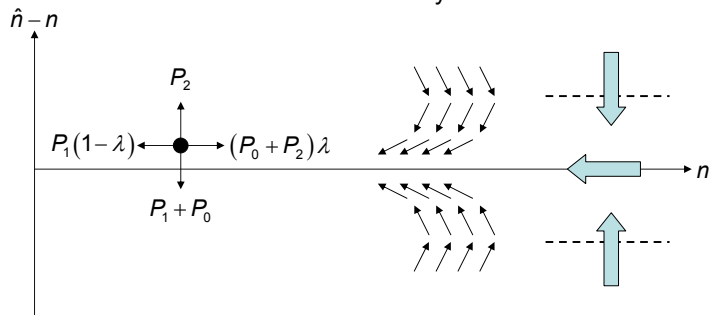
Adaptive slotted aloha

- Idea: Change the access probability with the number of stations
- How can we estimate the current number of stations in the system?
- Assume that stations can distinguish whether 0, 1, or more than 1 stations send in a time slot.
- Idea:
 - If you see that nobody sends, increase p .
 - If you see that more than one sends, decrease p .
- Model:
 - Number of stations that want to transmit: n .
 - Estimate of n : \hat{n}
 - Transmission probability: $p = 1/\hat{n}$
 - Arrival rate (new stations that want to transmit): λ ; note that $\lambda < 1/e$.



Adaptive slotted aloha 2

We have to show that the system stabilizes. Sketch:



$$\hat{n} \leftarrow \hat{n} + \lambda - 1, \text{ if success or idle}$$

$$\hat{n} \leftarrow \hat{n} + \lambda + \frac{1}{e-2}, \text{ if collision}$$



Adaptive slotted aloha Q&A

Q: What if we do not know λ , or λ is changing?

A: Use $\lambda = 1/e$, and the algorithm still works

Q: How do newly arriving stations know \hat{n} ?

A: We send \hat{n} with each transmission; new stations do not send before successfully receiving the first transmission.

Q: What if stations are not synchronized?

A: Aloha (non-slotted) is twice as bad

Q: Can stations really listen to all time slots (save energy by turning off)?

Q: Can stations really distinguish between 0, 1, and more than 1 sender?

A: No. One can use systems that only rely on acknowledgements...



Backoff Protocols

- Backoff protocols rely on acknowledgements only.
- Binary exponential backoff, for example, works as follows:
- If a packet has collided k times, we set $p = 2^{-k}$
Or alternatively: wait from random number of slots in $[1..2^k]$
- It has been shown that binary exponential backoff is not stable for any $\lambda > 0$ (if there are infinitely many potential stations)
[Proof sketch: with very small but positive probability you go to a bad situation with many waiting stations, and from there you get even worse with a potential function argument – sadly the proof is too intricate to be shown in this course ☺]
- Interestingly when there are only finite stations, binary exponential backoff becomes unstable with $\lambda > 0.568$; Polynomial backoff however, remains stable for any $\lambda < 1$.



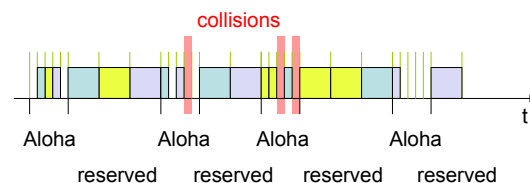
Demand Assigned Multiple Access (DAMA)

- Channel efficiency only 36% for Slotted Aloha, and even worse for Aloha or backoff protocols.
- Practical systems therefore use reservation whenever possible. But: Every scalable system needs an Aloha style component.
- Reservation:
 - a sender reserves a future time-slot
 - sending within this reserved time-slot is possible without collision
 - reservation also causes higher delays
 - typical scheme for satellite systems
- Examples for reservation algorithms:
 - Explicit Reservation (Reservation-ALOHA)
 - Implicit Reservation (PRMA)
 - Reservation-TDMA
 - Multiple Access with Collision Avoidance (MACA)



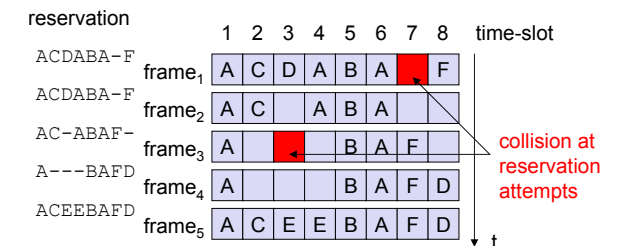
DAMA: Explicit Reservation

- *Aloha mode* for reservation: competition for small reservation slots, collisions possible
- *reserved mode* for data transmission within successful reserved slots (no collisions possible)
- it is important for all stations to keep the reservation list consistent at any point in time and, therefore, all stations have to synchronize from time to time



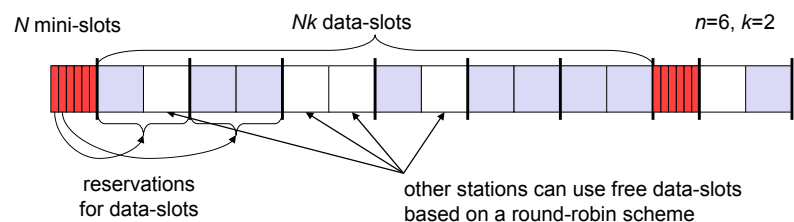
DAMA: Packet Reservation MA (PRMA)

- a certain number of slots form a frame, frames are repeated
- stations compete for empty slots according to the slotted aloha principle
- once a station reserves a slot successfully, this slot is automatically assigned to this station in all following frames as long as the station has data to send
- competition for this slots starts again as soon as the slot was empty in the last frame



DAMA: Reservation TDMA

- every frame consists of n mini-slots and x data-slots
- every station has its own mini-slot and can reserve up to k data-slots using this mini-slot (i.e. $x = nk$).
- other stations can send data in unused data-slots according to a round-robin sending scheme (best-effort traffic)



Multiple Access with Collision Avoidance (MACA)

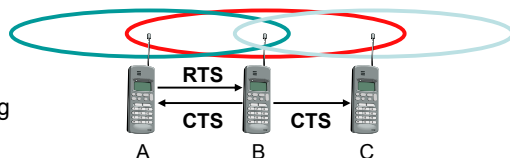
- Use short signaling packets for collision avoidance
 - Request (or ready) to send RTS: a sender requests the right to send from a receiver with a short RTS packet before it sends a data packet
 - Clear to send CTS: the receiver grants the right to send as soon as it is ready to receive
- Signaling packets contain
 - sender address
 - receiver address
 - packet size
- Example: Wireless LAN (802.11) as DFWMAC



MACA examples

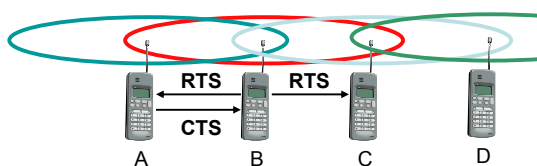
- MACA avoids the problem of hidden terminals

- A and C want to send to B
- A sends RTS first
- C waits after receiving CTS from B

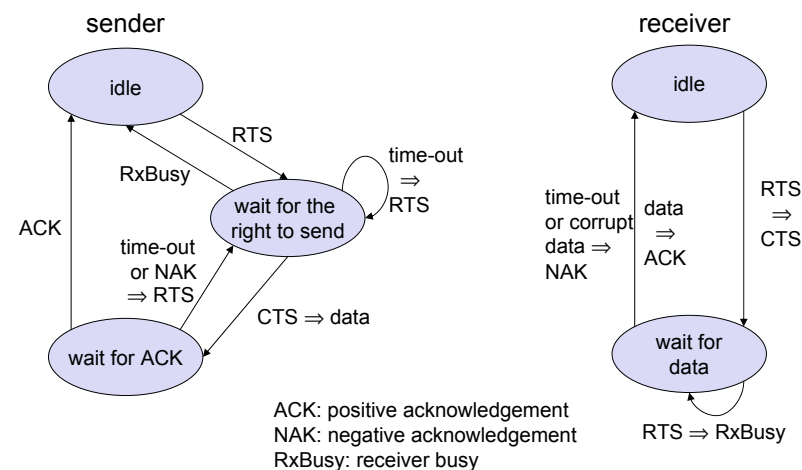


- MACA avoids the problem of exposed terminals

- B wants to send to A, and C to D
- now C does not have to wait for it cannot receive CTS from A



MACA variant: DFWMAC in IEEE802.11



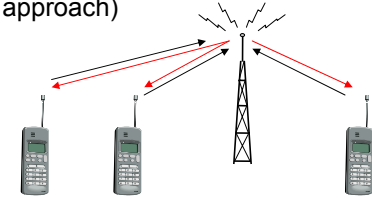
Polling mechanisms

- If one terminal can be heard by all others, this “central” terminal (a.k.a. base station) can poll all other terminals according to a certain scheme
 - Use a scheme known from fixed networks
- Example: Randomly Addressed Polling
 - base station signals readiness to all mobile terminals
 - terminals ready to send can now transmit a random number without collision with the help of CDMA or FDMA (the random number can be seen as dynamic address)
 - the base station now chooses one address for polling from the list of all random numbers (collision if two terminals choose the same address)
 - the base station acknowledges correct packets and continues polling the next terminal
 - this cycle starts again after polling all terminals of the list



Inhibit Sense Multiple Access (ISMA)

- Current state of the medium is signaled via a “busy tone”
- the base station signals on the downlink (base station to terminals) whether the medium is free
- terminals must not send if the medium is busy
- terminals can access the medium as soon as the busy tone stops
- the base station signals collisions and successful transmissions via the busy tone and acknowledgements, respectively (media access is not coordinated within this approach)
- Example: for CDPD (USA, integrated into AMPS)



Comparison SDMA/TDMA/FDMA/CDMA

Approach	SDMA	TDMA	FDMA	CDMA
Idea	segment space into cells/sectors	segment sending time into disjoint time-slots, demand driven or fixed patterns	segment the frequency band into disjoint sub-bands	spread the spectrum using orthogonal codes
Terminals	only one terminal can be active in one cell/one sector	all terminals are active for short periods of time on the same frequency	every terminal has its own frequency, uninterrupted	all terminals can be active at the same place at the same moment, uninterrupted
Signal separation	cell structure, directed antennas	synchronization in the time domain	filtering in the frequency domain	code plus special receivers
Advantages	very simple, increases capacity per km ²	established, fully digital, flexible	simple, established, robust	flexible, less frequency planning needed, soft handover
Dis-advantages	inflexible, antennas typically fixed	guard space needed (multipath propagation), synchronization difficult	inflexible, frequencies are a scarce resource	complex receivers, needs more complicated power control for senders
Comment	only in combination with TDMA, FDMA or CDMA useful	standard in fixed networks, together with FDMA/SDMA used in many mobile networks	typically combined with TDMA (frequency hopping patterns) and SDMA (frequency reuse)	still faces some problems, higher complexity, lowered expectations; will be integrated with TDMA/FDMA

J. Schiller

