

INTRODUCTION TO EECS II
**DIGITAL
COMMUNICATION
SYSTEMS**

**6.02 Fall 2012
Lecture #18**


- Shared medium → Media access (MAC) protocol
- Time division multiplexing
- Contention protocols: Aloha

6.02 Fall 2012
Lecture 18, Slide #1

Shared Media Networks



Satellite



ETHERnet



Wireless LANs






Cellular wireless

Individual images © source unknown. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <http://ocw.mit.edu/fairuse>.

6.02 Fall 2012
Lecture 18, Slide #2

Shared Communications Channels



- Basic idea in its simplest form: avoid collisions between transmitters – collision occurs if transmissions are concurrent
- Wanted: a communications protocol (“rules of engagement”) that ensures “good performance”
- Nodes may all hear each other perfectly, or not at all, or partially

6.02 Fall 2012
Lecture 18, Slide #3

“Good Performance”: What are the Metrics?

- High utilization
 - Channel capacity is a limited resource → use it efficiently
 - Ideal: use 100% of channel capacity in transmitting packets
 - Waste: idle periods, collisions, protocol overhead
- Fairness
 - Divide capacity equally among requesters
 - But not every node is requesting all the time...
- Bounded wait
 - An upper bound on the wait before successful transmission
 - Important for isochronous communications (e.g., voice/video)
- Dynamism and scalability
 - Accommodate changing number of nodes, ideally without changing implementation of any given node
- Not all protocols do well on all these metrics

6.02 Fall 2012
Lecture 18, Slide #4

Utilization

- Utilization measures the throughput of a channel:

$$U_{\text{channel}} = \frac{\text{total throughput over all nodes}}{\text{maximum data rate of channel}}$$

- Example: 10 Mbps channel, four nodes get throughputs of 1, 2, 2 and 3 Mbps. So utilization is $(1+2+2+3)/10 = 0.8$.
- $0 \leq U \leq 1$. Utilization can be less than 1 if
 - The nodes have packets to transmit (nodes with packets in their transmit queues are termed *backlogged*), but the protocol is inefficient.
 - There is insufficient *offered load*, i.e., there aren't enough packets to transmit to use the full capacity of the channel.
- With backlogged nodes, perfect utilization is easy: just let one node transmit all the time! But that wouldn't be fair...

6.02 Fall 2012

Lecture 18, Slide #5

Fairness

- Many plausible definitions. A standard recipe:
 - Measure throughput of nodes = x_i , over a given time interval
 - Say that a distribution with lower standard deviation is "fairer" than a distribution with higher standard deviation.
 - Given number of nodes, N , fairness F is defined as

$$F = \frac{\left(\sum_{i=1}^N x_i\right)^2}{N \sum_{i=1}^N x_i^2}$$

- $1/N \leq F \leq 1$, where $F=1/N$ implies single node gets all the throughput and $F=1$ implies perfect fairness.
- We'll see that there is often a tradeoff between fairness and utilization, i.e., fairness mechanisms often impose some overhead, reducing utilization

6.02 Fall 2012

Lecture 18, Slide #6

Channel Sharing Protocols

- Protocol \equiv "rules of engagement" for good performance
 - Known as media access control (MAC) or multiple access control
- Time division
 - Share time "slots" between requesters
 - Prearranged: time division multiple access (TDMA)
 - Not prearranged: contention protocols (e.g., Alohanet).
- Frequency division
 - Give each transmitter its own frequency, receivers choose "station"
 - Cf. lab for PS 6 – use different carrier frequencies & rcv filters
- Code division
 - Uses unique orthogonal pseudorandom code for each transmitter
 - Channel adds transmissions to create combined signal
 - Receiver listens to one "dimension" of combined signal using dot product of code with combined signal
 - Not covered in 6.02

6.02 Fall 2012

Lecture 18, Slide #7

Abstraction for Shared Medium

- Time is divided into *slots* of equal length
- Each node can start transmission of a packet only at the beginning of a time slot
- All packets are of the same size and hence take the same amount of time to transmit, equal to some integral multiple of time slots.
- If the transmissions of two or more nodes overlap, they are said to *collide* and none of the packets are received correctly. Note that even if the collision involves only part of the packet, the entire packet is assumed to be lost.
- Transmitting nodes can detect collisions, which usually means they'll retransmit that packet at some later time.
- Each node has a queue of packets awaiting transmission. A node with a non-empty queue is said to be *backlogged*.
- Depending on context, nodes may hear each other perfectly (eg, Ethernet), or not at all (e.g., satellite ground stations), or partially (e.g., WiFi devices or cell phones). For now, assume all nodes want to send packets to a fixed "master" (eg, base station)

6.02 Fall 2012

Lecture 18, Slide #8


Time Division Multiple Access (TDMA)

- Suppose that there is a centralized resource allocator and a way to ensure time synchronization between the nodes – for example, a cellular base station.
- For N nodes, give each node a unique index in the range $[0, N-1]$. Assume each slot is numbered starting at 0.
- Node i gets to transmit in time slot t if, and only if, $t \bmod N = i$. So a particular node transmits once every N time slots.
- No packet collisions! But unused time slots are “wasted”, lowering utilization. Poor when nodes send data in bursts or have different offered loads.

6.02 Fall 2012

Lecture 18, Slide #9

TDMA for GSM Phones



© Mozerati (top), source unknown (bottom). CC BY-SA. This content is excluded from our Creative Commons license. For more information, see <http://ocw.mit.edu/fairuse>.

6.02 Fall 2012

Lecture 18, Slide #10

The Aloha Protocol

Satellite clip art is in the public domain.

- Context: Norm Abramson, Hawaii
 - Developed scheme to connect islands via satellite network

Map © Europa Technologies, TerraMetrics, Google, and NASA. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <http://ocw.mit.edu/fairuse>.

6.02 Fall 2012


Lecture 18, Slide #11

Contention Protocols: Aloha (Simplest Example)

To improve performance when there are burst data patterns or skewed loads, use a contention protocol where allocation is not pre-determined.

Alohanet, designed by Norm Abramson et al. (Hawaii), was a satellite-based data network connecting computers on the Hawaiian islands. One frequency was used to send data to the satellite, which rebroadcast it on a different frequency to be received by all stations.

Stations could only “hear” the satellite, so had to decide independently when it was their turn to transmit.



Map © Europa Technologies, TerraMetrics, Google, and NASA. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <http://ocw.mit.edu/fairuse>.

6.02 Fall 2012

Lecture 18, Slide #12

Success, Idleness, Collisions

- Throughput = *Uncollided* packets per time interval
- Utilization of the above pic = Throughput / Channel Rate = 13/20 = .65

6.02 Fall 2012 Lecture 18, Slide #13

Slotted Aloha

- Aloha protocol – each node independently implements:
 - If a node is backlogged, it sends a packet in the next time slot with probability p .
- Assume (for now) each packet takes exactly one time slot to transmit (*slotted* Aloha)
- Utilization when N nodes are backlogged? The probability that exactly one node sends a packet.
 - prob(send a packet) = p
 - prob(don't send a packet) = $1-p$
 - prob(only one sender) = $p(1-p)^{N-1}$
 - There are $(N \text{ choose } 1) = N$ ways to choose the one sender

$$U_{\text{slotted Aloha}} = Np(1-p)^{N-1}$$

6.02 Fall 2012 Lecture 18, Slide #14

Maximizing Utilization

Utilization for slotted Aloha (N=10)

to determine maximum: set $dU/dp = 0$, solve for p .

Result: $p = 1/N$, so

$$U_{\text{max}} = \left(1 - \frac{1}{N}\right)^{N-1}$$

As $N \rightarrow \infty$, $U_{\text{max}} \rightarrow \frac{1}{e} \approx 37\%$

$$\begin{aligned} \ln\left(1 - \frac{1}{N}\right)^{N-1} &= (N-1)\ln\left(1 - \frac{1}{N}\right) \\ &= (N-1)\left(-\frac{1}{N} - \frac{1}{2N^2} - \frac{1}{3N^3} - \dots\right) \\ &= -1 + \frac{1}{N} + \frac{1}{6N^2} + \frac{1}{12N^3} + \dots \\ &= -1 \text{ as } N \rightarrow \infty \end{aligned}$$

6.02 Fall 2012 Lecture 18, Slide #15

Simulation of Slotted Aloha (N=10)

Utilization = .38, Fairness = .98

6.02 Fall 2012 Lecture 18, Slide #16

Stabilization: Selecting the Right p

- Setting $p = 1/N$ maximizes utilization, where N is the number of *backlogged* nodes.
- With bursty traffic or nodes with unequal offered loads (aka *skewed loads*), the number of backlogged is constantly varying.
- Issue: how to dynamically adjust p to achieve maximum utilization?
 - Detect collisions by listening, or by missing acknowledgement
 - Each node maintains its own estimate of p
 - If collision detected, too much traffic, so decrease local p
 - If success, maybe more traffic possible, so increase local p
- “Stabilization” is, in general, the process of ensuring that a system is operating at, or near, a desired operating point.
 - Stabilizing Aloha: finding a p that maximizes utilization as loading changes.

6.02 Fall 2012

Lecture 18, Slide #17


Binary Exponential Backoff

- Decreasing p on collision
 - Estimate of N (# of backlogged nodes) too low, p too high
 - To quickly find correct value use multiplicative decrease: $p \leftarrow p/2$
 - k collisions in a row: p decreased by factor of 2^{-k}
 - Binary: 2, exponential: k , back-off: smaller $p \rightarrow$ more time between tries
- Increasing p on success
 - While we were waiting to send, other nodes may have emptied their queues, reducing their offered load.
 - If increase is too small, slots may go idle
 - Try multiplicative increase: $p \leftarrow \min(2*p, 1)$
 - Or maybe just: $p \leftarrow 1$ to ensure no slots go idle

6.02 Fall 2012

Lecture 18, Slide #18


Simulation of Stabilized Aloha



6.02 Fall 2012

Lecture 18, Slide #19

Node probabilities over time (different run from previous page)




6.02 Fall 2012

Lecture 18, Slide #20


What Went Wrong?

- Starvation
 - Too many successive failures $\rightarrow p$ very small \rightarrow no xmit attempts
 - Result: significant long-term unfairness
 - Try a reduction rule with a lower bound: $p \leftarrow \max(p_{min}, p/2)$
 - Choosing $p_{min} \ll 1/\max(N)$ seems to work best
- But there's another problem: capture effect




Limiting the Capture Effect


- Capture effect
 - A successful node maintains a high p (avg. near 1)
 - Starves out other nodes for short periods
 - Try an increase rule with an upper bound: $p \leftarrow \min(p_{max}, 2*p)$



Node Probabilities: pmax=.25, pmin=.01 (different run from previous page)



Aloha in Pictures: Collisions



Collisions! Packet is 5 slots long in this picture

- A collision occurs when multiple transmissions overlap in time (even partially)
- Throughput = *Uncollided* packets per second
- Utilization = Throughput / Channel Rate

Dear Professor Balakrishnan,

I just wanted to let you know how my IAP has been going **because I've been using a few 6.02 concepts**. I'm interning at Quizlet.

... I also worked on making a testing tool. I set up a node server to manage many headless browsers (using phantomjs). I realized that phantomjs wasn't built to manage the number of connections I required. **Then I noticed that this was similar to the problem of connecting multiple clients to the same router**. So I set up a system similar to one of the 6.02 labs in which everyone was assigned a random time to start their connection, over a period of time. When a browser started initialising, I put on a lock that made other browsers attempting to connect wait another random time period. Once the "messy" part of the initialisation was done, I unset the lock in order to allow other clients to connect.

I had to implement various other 6.02ish features like exponential back off, and I've also noticed that node.js is (of course) a very 6.02 type of project. For example, it has a concept of heartbeats.

Best, Chase (Fall 2011 student that sat in the front row and whom you approached at the gym)

6.02 Fall 2012

Lecture 18, Slide #25

Spot Quiz

- In Aloha, each node maintains a variable, p . What does p represent? [P(x) is "probability of event x"]
 - P(node being backlogged)
 - P(backlogged node sends in a timeslot)
 - P(packet transmission is received correctly)
 - P(time slot is kept idle)
- In **stabilized** Aloha, the value of p never goes below p_{min} . Why should p_{min} not be too small?
 - To increase the utilization
 - To avoid extreme unfairness
 - To reduce the problems caused by the "capture effect"
 - To reduce the number of collisions
- Slotted Aloha, packet size 1 slot, N backlogged nodes, each node has a **fixed** p . Calculate: P(**collision** in a timeslot)

6.02 Fall 2012

Lecture 18, Slide #26

Spot Quiz

- In Aloha, each node maintains a variable, p . What does p represent? [P(x) is "probability of event x"]
 - P(node being backlogged)
 - B. P(backlogged node sends in a timeslot)
 - P(packet transmission is received correctly)
 - P(time slot is kept idle)
- In **stabilized** Aloha, the value of p never goes below p_{min} . Why should p_{min} not be too small?
 - To increase the utilization
 - B. To avoid extreme unfairness
 - To reduce the problems caused by the "capture effect"
 - To reduce the number of collisions
- Slotted Aloha, packet size 1 slot, N backlogged nodes, each node has a **fixed** p . Calculate: P(**collision** in a timeslot)
 Soln: P(collision) = P(≥ 2 xmits) = $1 - P(\text{no xmit}) - P(1 \text{ xmit}) =$
 $P(\text{no xmit}) = (1-p)^N$; $P(1 \text{ xmit}) = Np(1-p)^{(N-1)}$
 Therefore, P(collision) = $1 - (1-p)^N - Np(1-p)^{(N-1)}$.

6.02 Fall 2012

Lecture 18, Slide #27

MIT OpenCourseWare
<http://ocw.mit.edu>

6.02 Introduction to EECS II: Digital Communication Systems
Fall 2012

For information about citing these materials or our Terms of Use, visit: <http://ocw.mit.edu/terms>.