Exercise Sensor Networks – (till may 23, 2005)

Lecture 4: MAC and energy efficiency

Exercise 4.1: Genie Aided Aloha

Genie-aided Aloha was an estimate for the energy efficiency of the Aloha protocol. Is GAA better than pure Aloha in every case and if not when and why?

Solution:

With an increasing arrival rate the energy consumption of GAA is converging against the one of pure Aloha because empty frames (with no arrival) become increasingly rare. In general all approaches trying to avoid idle listening lose their effectiveness if a lot of communication is going on.

Exercise 4.2: Slotted Aloha

In what way does slotted Aloha differ from pure Aloha with regard to the channel access? Try to quantify how the two approaches differ (in this context the packet delivery rate is not important).

Solution:

Pure Aloha allows to access the medium at any time (apart from the question whether the packet will survive). In Slotted Aloha a station has to wait for the beginning of the next frame. The average waiting time will be ½ frame time. If the channel is idle this waiting time is unnecessary. So if usually only one station is sending pure Aloha is better suited than Slotted Aloha.

©Thomas Haenselmann – Applied Computer Science IV, University of Mannheim

Exercise Sensor Networks

Lecture 4: MAC and energy efficiency

Exercise 4.3: Comparison of medium access approaches

Why is the delivery rate of 1-persistent CSMA better than the one of Slotted Aloha and why do both of them converge against the same delivery rate in very busy channels?

Solution:

As 1-persistent Aloha uses “carrier sense”, it listens to the channel and tries to access it only after the current transmission has finished. In case of high packet rates the medium is almost always occupied. All emerging transmission desires will be queued by the station until the medium is clear again. Then all waiting stations will greedily access the channel and their packets collide. Exactly the same is true for Slotted Aloha. Actually the only difference is the reason why stations wait. In Slotted Aloha waiting is controlled by the clock, in all CSMA approaches like 1-persistent Aloha it is triggered by the carrier sense. But both are greedy when it come to accessing the channel so the delivery rate will converge against zero with an increasing arrival rate.

In case of low arrival rates on the channel 1-persistent Aloha has the advantage that is can access the channel immediately, whereas in Slotted Aloha ½ frame will pass on average. This does not only increase the delay but it is at the same time a waste of channel capacity.

©Thomas Haenselmann – Applied Computer Science IV, University of Mannheim

Exercise Sensor Networks

Lecture 4: MAC and energy efficiency

Exercise 4.4: p-persistent CSMA

1-persistent CSMA means to send instantly at the beginning of a frame time while non-persistent CSMA means that a random time has to pass before sending in case of an occupied channel. In between those extremes a probability p can be chosen which is the likeliness for sending at the beginning of the next frame time.

For what reason may probabilities in between 0 and 1 be more optimal than than 0 or 1?

Solution:

In the case of 1-persistent CSMA packets destroy themselves as stated in 4.3 because nodes access the channel greedily.

0-persistence means the a station will never access the medium right after the current station has finished its transmission. This does on one hand mitigate the problem of collisions making them highly unlikely but it also waists the time between the end of an ongoing packet and the next attempt to access the channel. Theoretically the chance that a station will access the channel right after an ongoing packet is zero.

So if a small number of stations were greedy anyhow they could use the free capacity which nobody else takes as a result of an exaggerated politeness. The probability that a station will access the channel after the transmission without waiting a random time should be high enough that it actually happens (given a certain number of stations) but it should be low enough that the probability of a collision is still low.

As a consequence the total number of stations using a common channel will determine the choice of p.

©Thomas Haenselmann – Applied Computer Science IV, University of Mannheim

Exercise Sensor Networks

Lecture 4: MAC and energy efficiency

Exercise 4.5: Aloha with Preamble Sampling

Basic consumption	: 8 mA	$b^{PAS} = 1 - e^{-gN(T_P+T_S+T_R+T_A)}$	P' for incoming message
Energy f. sending	: 12 mA	$b_1^{PAS} = 1 - e^{-g(T_P+T_S+T_R+T_A)}$	P' for sending a message
Energy f. receiving	: 6 mA	$Pow^{PAS} = b_1^{PAS} P_{TX} + (b^{PAS} - b_1^{PAS}) P_{RX}$	Mean consumption f. send. and recv.
Sleep mode	: ~0mA		

Let the length of a packet T_M be 0,8 times the frame time, and let the time T_S for switching the transceiver between sending and receiving and the time T_A for an acknowledgment be 0.1 times the frame time, so that a full transmission attempt occupies exactly one frame time (0.8+0.1+0.1). The preamble T_P for waking up the neighbors should take another full frame time which means that a sensor node has to wake up once per frame time. The sending rate g should be 0.01 (attempts to send per frame time) and the total number of nodes should be 10.

a) How high is the mean energy consumption in this scenario?
b) How high is the energy consumption if waking up and listening to the channel consumes T_W=14,0mA of energy. In order to check, whether the channel is free or occupied a node has to stay awake and keep listening for at least 1% of the frame time.

©Thomas Haenselmann – Applied Computer Science IV, University of Mannheim

Exercise Sensor Networks

Lecture 4: MAC and energy efficiency

Solution:

a) How high is the mean energy consumption in this scenario?

Basic consumption : 8 mA
Energy f. sending : 12 mA
Energy f. receiving : 6 mA
Sleep mode : ~0mA

$$b^{PAS} = 1 - e^{-0,01 \times 10(1+0,8+0,1+0,1)} \approx 0,181$$

$$b_1^{PAS} = 1 - e^{-0,01(1,0+0,8+0,1+0,1)} \approx 0,02$$

$$Pow^{PAS} = 0,02(12+8) + (0,181 - 0,02)(6+8) = 2,654$$

b) How high is the energy consumption if waking up and listening to the channel consumes T_W=14,0mA of energy. In order to check, whether the channel is free or occupied a node has to stay awake and keep listening for at least 1% of the frame time.

$$Pow_{+Wakeup}^{PAS} = 2,654 + (1 - 0,181 - 0,02)0,01 \times 14 = 2,654 + 0,112 = 2,766$$

©Thomas Haenselmann – Applied Computer Science IV, University of Mannheim

Exercise Sensor Networks

Lecture 4: MAC and energy efficiency

Exercise 4.6: Optimizing the preamble

In the following scenario there are 100 nodes. Each of them produces and arrival rate of 0.00005 packets per frame (so communication is somewhat rare). A packet occupies a full frame length. Acknowledgments are not implemented, at least not on the mac layer. Nodes that encounter an active channel stay awake until the channel is clear again. Transmitting a packet is as expensive as receiving one.

a) Imagine to be a node. How many percent of our lifetime do we spend to send a packet, do we spend to hear a packet and do we sleep?

$$b_{snd} = 1 - e^{-0.00005}$$
$$b_{rcv} = 1 - e^{-0.00005 \times 100} - b_{snd} = 1 - e^{-0.00005 \times 100} - b_{snd} = 1 - e^{-0.00005 \times 99} \times e^{-0.00005} - b_{snd} = e^{-0.00005}(1 - e^{-0.00005 \times 99})$$
$$b_{idl} = 1 - b_{rcv} - b_{snd} = 1 - (1 - e^{-0.00005 \times 100} - b_{snd}) - b_{snd} = e^{-0.00005 \times 100}$$

b) A node wakes up during a transmission. How many percent of the preamble and the actual packet does the node hear on average?

Solution:

The message consists of the preamble PA (measured in frame times) and a message so the full length of a transmission is (PA+1). The node can wake up at any time, e.g., in the beginning of the transmission of in the end. So it hears 0.5x(PA+1) on average.

©Thomas Haenselmann – Applied Computer Science IV, University of Mannheim

Exercise Sensor Networks

Lecture 4: MAC and energy efficiency

Exercise 4.6: Optimizing the preamble

In the following scenario there are 100 nodes. Each of them produces and arrival rate of 0.00005 packets per frame (so communication is somewhat rare). A packet occupies a full frame length. Acknowledgments are not implemented, at least not on the mac layer. Nodes that encounter an active channel stay awake until the channel is clear again. Transmitting a packet is as expensive as receiving one.

c) Aloha with preamble sampling requires to wake up once within the time of a full preamble. We assume that waking up, listening into the channel and going pack to sleep in case of silence is expensive and costs 10% of the energy require to transmit a packet. What is the optimal length for the preamble (measured in frame times) in this situation?

If the preamble is PA-times a frame time a node wakes up 1/PA times per frame time. Within this time it needs 0.1 time the energy for listening to a full frame so the total amount if 0.1/PA. The probability that this happens is b_{idl}.

$$P = b_{idl} \times 0.1/PA$$

Next the energy needed to send a packet (and a full preamble is added):

$$P = b_{idl} \times 0.1/PA + b_{snd}(PA+1)$$

Finally the probability to encounter an ongoing preamble (PA/2) and the following data packet is added:

$$P = b_{idl} \times 0.1/PA + b_{snd}(PA+1) + b_{rcv}(PA/2+1)$$

©Thomas Haenselmann – Applied Computer Science IV, University of Mannheim

Exercise Sensor Networks

Lecture 4: MAC and energy efficiency

Exercise 4.6: Optimizing the preamble

c) continued:

The graph of the energy-requirement function depending on the preamble length looks like that:

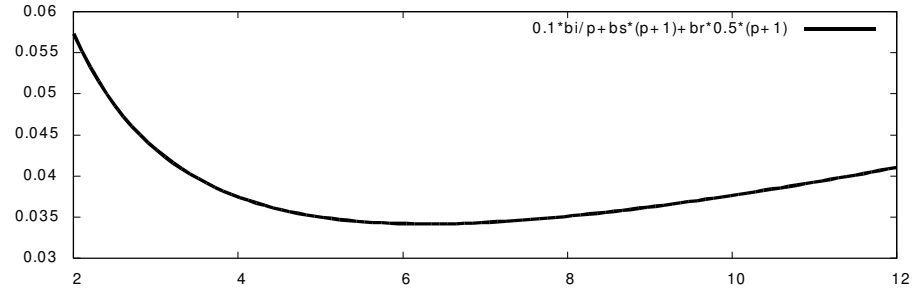

for what PA is P minimal?

$$P = b_{idl} \times 0.1/PA + b_{snd}(PA+1) + b_{rcv}(PA/2+1)$$

$$P' = -0.1\, b_{idl} PA^{-2} + b_{snd} + 0.5\, b_{rcv} = 0$$

$$0.1\, b_{idl} PA^{-2} = b_{snd} + 0.5\, b_{rcv} \Leftrightarrow PA = \pm\sqrt{\frac{b_i}{10\, b_s + 5\, b_r}} = \pm\sqrt{\frac{e^{-0.00005\times 100}}{10(1-e^{-0.00005}) + 5e^{-0.00005}(1-e^{-0.00005\times 99})}} \approx 6.29$$

©Thomas Haenselmann – Applied Computer Science IV, University of Mannheim

Exercise Sensor Networks

Lecture 4: MAC and energy efficiency

Exercise 4.7: Simulation of the Aloha approach

a) Compile the program on our computer. You may want to port it to Java (only marginal changes are necessary).

b) Change the program in order to simulate slotted Aloha. This can be accomplished by swapping few lines of code and adjusting the arrival rate.

Solution:

The only difference is that the stations are not triggered in each simulation time unit but only once per frame (this is accomplished by ((time % FRAME_LENGTH) == 0). However, the probability of triggering a station has to rise by the factor 100 since we trigger 100 times less often.

```
for(long time = 0; time < MAX_TIME; time++) {

// Trigger stations only at beginning of frame time

if((time % FRAME_LENGTH) == 0)
  for(long station_index = 0; station_index < NO_STATIONS; station_index++)
    if((abs(rand()) % (NO_STATIONS*100)) < arrival_rate) station[station_index].TriggerSend(time);

  if(survival_timer != -1) survival_timer++;
  if(survival_timer == FRAME_LENGTH) {
    successful_packets++;
    survival_timer = -1;
  } // if
} // for
```

See our homepage for the full code.

©Thomas Haenselmann – Applied Computer Science IV, University of Mannheim

Exercise Sensor Networks

Lecture 4: MAC and energy efficiency

Exercise 4.6: Simulation of the Aloha approach

c) Now extent the class 'Station' in order to simulate p-persistent Aloha. It should be possible to vary p between 0 (for non-persistent aloha) and 1 (for 1-persistent Aloha). Which instance of p is optimal for 100 stations?

Solution:

See our homepage for the full code.

©Thomas Haenselmann – Applied Computer Science IV, University of Mannheim