Lab 04: NS2 based study of the performance of CSMA/CD (802.3)

OSL, Monday 27 Aug, 2012

Objective:

Understand the performance of CSMA/CA in terms of system efficiency and fairness measure.

General instructions:

1. This lab is to be done **individually**.
2. Download the file "cs378-lab04-student-files.tgz" and unzip it. You will find all relevant files needed for this lab in this directory.
3. Create a directory called <rollnumber>_lab04. As you proceed with the lab instructions below, note down observations or relevant output from whatever you do in a file named "lab04.txt" using a text editor.
4. Also add to this directory any written code along with output files. You will find more details of this in the specific exercises.

Lab Instructions:

Exercise 1: CSMA Warmup

For this exercise, you will be using the file named cs378-lab04-ns-csmacd.tcl to run the experiments. This script essentially creates a LAN based on 802.3(CSMA/CD), which is the name of the MAC standard for Ethernet. It takes as inputs the number of nodes that will be part of the LAN and a seed for the random number generator. Each node in the LAN generates CBR traffic at the rate of 100kbps. Carefully go through the script file and answer the following in lab04.txt.

1. Which command was used to create the LAN? What are its arguments? How was the MAC protocol specified?
2. What is the data rate (bandwidth) of the LAN?
3. What is the packet size (bytes) used in the generated traffic? What is the interpacket gap?
4. What is the duration of the simulation when **traffic is active**?

Execute the script with number of nodes as 4 and seed as "1234". Basically: "ns ns2-csmacd-lab04.tcl -nn 4 -seed 1234". Look at the NAM output. You will see a total of 6 nodes, even though you set only 4. Node 0 in this exercise is a sink node i.e it does not generate any traffic and just sinks traffic generated by other nodes. Nodes of relevance to us are nodes 1,2,3,4. Node 5 is not listed in NAM but if you look at the output file "csmacd.tr", you will see node 5 listed. The LAN bus is modelled in ns2 as a node instead of as a link, where any node wishing to transmit a packet on the link, essentially gives it to node 5. You will see this operation listed by a "h" in the output file. Node 6 is again a dummy node, created just to get the orientation right for easy visualization in NAM. You can ignore this node and the corresponding link between it and node 0. (For added fun, you can comment the code corresponding to node6, to see how Nam orients the nodes without it.)

Sample Run: Let us process the results for one sample run before we perform other runs. Look at the csmacd.tr file generated above using nodes as 4 and seed as 1234. Answer the following in

lab04.txt.

1. What are the total number of packets (P) succesfully received by the reciever?
2. What are the total number of packets dropped (at different senders) during the run? Why are these packets dropped? What is the **average** packet drop percentage at a sender?
3. What is the system throughput and efficiency? System throughput is defined as the total number of bits succesfully received (P * packet size * 8) divided by the simulation duration. System efficiency is the ratio of system throughput by the data rate of the LAN. Express system efficiency as a percentage. Why is the system efficiency not 100%?
4. What are the individual throughputs of nodes 1,2,3 and 4? For this, you need to calculate the number of packet received succesfully at node 0 from nodes 1,2,3 and 4 respectively. So, a total of 4 values, one corresponding to each node. Check that the sum of individual throughputs sums to the system throughput calculated in Q.3
5. Trace the collision profile for the first packet (sequence number 0) generated by nodes 1,2,3 and 4? In other words, what is the time at which these packets got generated? When did they get transmitted? Did they collide? If so, who won the collision etc.... You just need to look till time 0.54 sec in the csmacd.tr to figure this out.

Exercise 2: Throughput, Efficiency and Packet Drops

Now we will run experiments as we vary the number of nodes. You will need to automate this process using a for loop in bash. **Important: For these runs, specify the seed as "last four digits of your roll number". Also comment out the line "exec nam out.nam", otherwise you will have to close a lot of nam windows.**

1. Vary the number of nodes from 1 to 20 i.e. nn takes on values 1,2 ....., 20 and for each run specify the seed as calculated above based on your roll number. This gives a total of 20 runs.
2. For each run, calculate the system efficiency, system throughput and total drops. Output this to a file overall-stats.txt, where the first column is the number of nodes, the second column is the system throughput, the third column is the system efficiency and the fourth column is the total drops. In summary after all the runs, the file should have 20 rows and 4 columns, one row generated from each run.
3. Use gnuplot to generate three graphs. One plots system throughput as a function of the number of nodes ("my-thr.eps"), the second plots the system efficiency (my-eff.eps") as a function of nodes. The third plots total drops as a function of nodes (my-drops.eps). Look at the samples "thr-nodes.eps" , "eff-nodes.eps" and "drop-nodes.eps"
4. Comment on the graphs plotted in lab04.txt. What is the average packet drop percentage experienced by a sender, when number of nodes in the network is 10 vs 20.
5. Include your bash scripts, overall-stats.txt, my-thr.eps, my-eff.eps, my-drops.eps in the submission directory. **DO NOT** submit the trace files for these runs since they will be quite a few.

Exercise 3: Faireness Run

Run the experiment by setting the number of nodes as 20 and the seed corresponding to your roll number. For this run, calculate the individual throughputs for each of the individual nodes i.e. 1,2,3.....20. Output this to a file "run20-node-stats.txt", where the first column is the node id and

the second column is the throughput obtained by that individual node. This file will have 20 rows and 2 columns. Use gnuplot to plot the individual node throughputs (“my-fairness.eps”). Look at the sample “fairness.eps”.

1. Comment on the graphs obtained.
2. Use the Jain's fairness index (http://en.wikipedia.org/wiki/Fairness_measure) to comment on the fairness of Ethernet.
3. Submit the trace file csmacd.tr, run20-node-stats.txt and my-fairness.eps

Submission instructions

The directory named <rollnumber>_lab04 that you will submit should contain the following files:

1. lab04.txt
2. bash script
3. overall-stats.txt
4. my-thr.eps
5. my-eff.eps
6. my-drops.eps
7. csmacd.tr (corresponding to the run with 20 nodes and roll number seed)
8. run20-node-stats.txt
9. my-fairness.eps

Now tar it as follows:
tar -zcvf <rollnumber>_lab04.tgz <rollnumber>_lab04/

Submit the file <rollnumber>_lab04.tgz via moodle for grading.