TSIN02 Internetworking

Exercise class 5 problems

Exercise 1: You have started a data center provider company. In one part of a server hall, you have 1000 helper (H) servers that run helper functions to be called by the 1000 main computational servers (M). You must connect these two groups of computers - good connectivity is important! You choose to connect the M- and H-servers through a Clos network with the M-servers on one side, and the H-servers on the other side.

The Clos network has 2 x 3-type input switches, which means that the output switches are 3 x 2-type switches.

a) How many middle switches do you need?
b) How many connectors on each side do the middle switches have?
c) Is your Clos network non-blocking or re-arrangeably non-blocking? Motivate your answer!
d) We know that big switches are expensive. Replace each middle switch with a group of smaller switches! Choose a Clos network for this task, and maintain the non-blocking or re-arrangeably non-blocking property above. Describe the used switch types, i.e., the number of input and output connectors, for each Clos network switch.

Exercise 2: A cloud provider wants good connectivity in its data center. Therefore, a fat tree is used instead of a standard tree structure. You are to investigate how to build this with very simple switches.

Sketch a fat tree built only with switches with the following specification: each switch has 2 unidirectional inputs and 2 unidirectional outputs. The fat tree should have 8 bidirectional inputs in total, be re-arrangeably non-blocking, and built from a 5-stage Clos network. Motivate your network building in a systematic way by starting with an unfolded 3-stage Clos network, modify step by step your network while commenting on the re-arrangeably non-blocking property, and fold in the last step. Make drawings of the network as you develop it.

Exercise 3: You need to connect 8 helper (H) servers that run helper functions to be called by the 8 main computational servers (M). You choose to connect the M- and H- servers through a re-arrangeably non-blocking three-stage Clos network with the M- servers on one side, and the H- servers on the other side. What size and numbers of input, middle, and output switches can you use to build a Clos network that requires a minimum number of total switches?

Exercise 4: You need to connect 6 helper (H) servers that run helper functions to be called by the 6 main computational servers (M). You choose to connect the M- and H- servers through a Clos network with the M- servers on one side, and the H-servers on the other side. The available switches are of size 2 x 2 and 3 x 3, which can be used to build the Clos network. The price of a 2 x 2 and a 3 x 3 switch is 150 SEK and 300 SEK, respectively.

a) Is it possible to build a strict-sense non-blocking Clos network from the available two sizes switches? Motivate your answer!
b) Motivate why it is possible to build a re-arrangeably non-blocking Clos network from the available two sizes switches.
c) You have a limited budget and need to build a cost-efficient re-arrangeably nonblocking Clos network. What size and numbers of input, middle, and output switches you can use to build a cost-efficient Clos network? What will be the total cost of your network? Neglect the cost of the connectors.

Exercise 5: A research center in Stockholm produces 600 GB of new data for every wet lab experiment. Assume the data generated can be easily parallelized, with a negligible overhead.
a) Suppose the Amazon Web Services (AWS) sells CPU hours at the price of 0.10 SEK/hr per Elastic Compute Cloud (EC2) instance, where each instance takes 2 hours to process 1 GB of the experimental data. The data transfer fee is 0.15 SEK/GB. What is the price the research center will need to pay for processing the experiment using the EC2 service?
b) Suppose the data transfer rate from the research center to AWS is 20 Mbps. What is the total time required to transmit and process the experiment data using the EC2 service?
c) The research center has 24 computers itself, each taking 2 hours to process a GB of data. Suppose the maintenance fee (including electricity, software, hardware, etc) is 15 SEK per computer per experiment. What is the total amount of time and cost required to process the experiment? Will the research center be willing to use the EC2 service?