CS 653 Mobile Computing *Spring 2014*

Problem Set 3 Solutions

Problem 1 [1 mark]

Consider a multihop chain of N nodes, 1, 2, …, N. Every node is within radio range of its immediately adjacent nodes on the chain, and data is routed and forwarded along the chain from node 1 to node N. Every node has a single half duplex radio that can only send or receive (but not both) on a single channel. Every node interferes with nodes up to F hops away. What is the maximum long term data transfer throughput (in packets per second) that this chain of nodes can sustain, under the following conditions:

a. The physical layer throughput of any link in the chain is R packets per second.
b. The physical layer throughput of the link between nodes i and i+1 is R/i packets per second, for i = 1 to N-1.

Your answers should be in terms of N, F, and R.

Solution

When a certain node is transmitting, a few other nodes in the vicinity will not be able to transmit due to constraints such as half-duplex radios, interference etc. Let us call this set of nodes the interfering set. Beyond the interfering set, the rest of the transmissions can happen in parallel. From the reasoning similar to the flush paper, the length of the interfering set (in terms of hops) is min (N-1, F+2). Let is denote this by the variable "X". Note that with N nodes, we have a chain of N-1 hops, hence the slight difference with the formula in the Flush paper.

a. Only one of X nodes will be able to send at a time. When the first node sends R packets in one unit of time, these packets take a time of X units to get through the pipeline of X hops. So maximum throughput will be R/X, where X = min(N-1,F+2) as described above.
b. Since the physical layer throughputs are progressively decreasing along the chain, the maximum throughput is that which can be sustained by the last set of X hops. When the first node sends R packets in one unit of time, the node "i" takes R/i units of time to send these R packets. The amount of time taken by the last X hops in the pipeline to transmit R packets is given by (N-1) + … + (N-X). Let us denote this time by T. In this time T, all other set of X nodes can transmit R packets, and therefore will not be the bottleneck. Therefore the throughput of the system, which is the throughput of the last X hops is R/T.
   Note that many students have reasoned that the bottleneck rate is R/(N-1), so the maximum throughput is R/((N-1)*X). This reasoning is incorrect, and is based on the assumption that each of the nodes in the last interfering set of X nodes will take a time of (N-1) units to transmit R packets. This overestimates the time taken, hence underestimates the throughout. The correct answer above gives a larger value of throughput than this incorrect answer.

Problem 2 [2 marks]

In which of the following situations is Mobile IP useful in maintaining network connectivity during mobility? If useful, please explain **how** Mobile IP is useful in the said situation.

a. A client who only initiates TCP connections moves from one location to another and obtains a new IP address. It does not have any outstanding TCP connections during the move.
b. A client who only initiates TCP connections moves from one location to another and obtains a new IP address. It has an ongoing TCP connection during the move that it does not want to disrupt.
c. A server that accepts TCP connections moves from one location to another and obtains a new IP address. It updates it's DNS entries very quickly (with almost zero delay, during which no connections arrive at the server) after the move.
d. A server that accepts TCP connections moves from one location to another and obtains a new IP address. However, updating DNS to reflect the new IP address takes a long time, during which the server may potentially receive incoming connections.

Solution

Mobile IP is useful in cases (b) and (d) above. In cases (b) and (d) the tunneling mechanism of Mobile IP is useful to maintain ongoing connections or start new connections at the old IP address, while the mobile node has moved to the new IP address. In case (a), the client can initiate a new connection with new IP address, and no packet will arrive at the old IP. Hence none of the Mobile IP mechanisms will kick in. In case (c), assuming the server has no ongoing connections, and all clients resolve DNS name first, the DNS update will ensure that new connections arrive at the new IP address. However, if clients use IP address directly, some packets may arrive at the old IP of the server, and Mobile IP may be needed. So the answer to part (c) depends on your assumptions.

Problem 3 [2 marks]

Consider the following multihop topology of 8 nodes, A through H.

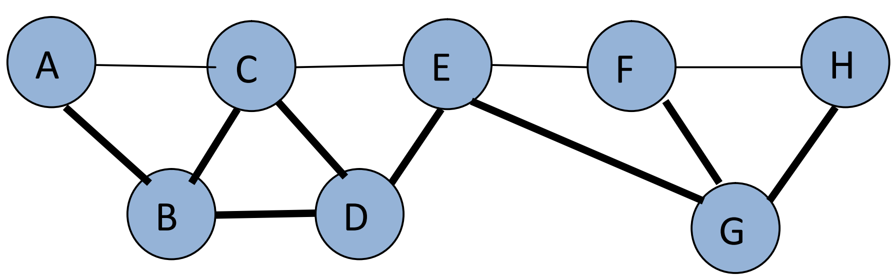

The links shown with then lines have a metric 1 and the links with thick lines have a metric 3. Nodes that do not have a link between them are not within radio range of each other. The nodes use a DSDV like routing protocol. For the sake of simplicity, assume that H is the only possible destination in the routing and forwarding tables. For this destination H, each node maintains a set of routes, which are tuples of the form (next-hop, metric). The node then picks one of these routes as its best route, and periodically advertises it to neighbors using DSDV-like routing messages. For each case below**, list all known routes (next-hop, metric) to the destination H at nodes A and E.** Also indicate which of these is the **best route** used by the node for forwarding. You may provide the answers assuming the routing and forwarding information has converged to a steady state after any changes.

a. When the topology is static as shown in the figure above.
b. When the links C-E and E-F go down.
c. When nodes C and F go down.
d. When node E goes down.

Solution

Every node will potentially learn of multiple routes in its routing table, one from each of its neighbors that has a route. Out of them, it will pick the route with the smallest metric as its best route and use it for forwarding. Every node also periodically advertises its best route and its metric to all its neighbors. The best route is marked with a * below. You are required to correctly identify all the visible routes in addition to the best path.

a. At E the routes are (F, 2)*, (G, 6), (C, 3), (D, 5). At A the routes are (C, 4)*, (B, 9).
b. E: (G, 6)*, (D, 9). At A: (C, 13)*, (B, 15)
c. E: (G, 6)*, (D, 9). At A: (B, 15)*
d. A has no valid route to H