ECSE-6961:Internet Protocols
Quiz 2

Time: 60 min (strictly enforced)

Points: 50 { 5 questions, each 10 points}

YOUR NAME:

Be brief, but DO NOT omit necessary detail

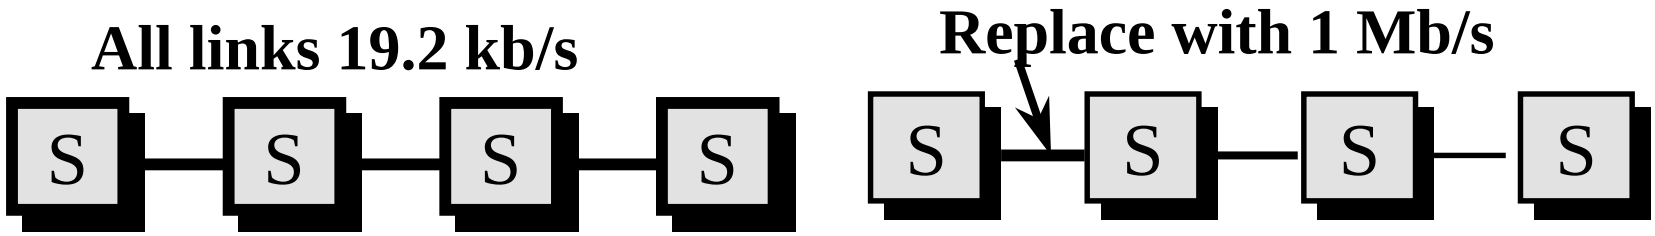

File Transfer time = ***5 mins***

File Transfer Time = ***7 hours***

- **1. Explain the problem in the figure to the right. If any other link were replaced by the 1 Mbps link, would you still encounter the congestion problem ? Why ? What does this tell you about placement (provisioning) of high speed links in a network ?**

- **2. We would like to control the rates of TCP flows using a network component which can manipulate TCP acks. You are given the desired rate (R) and round trip time (RTT) for a TCP flow. How would you enforce this maximum rate and spread packet transmissions over an RTT by only manipulating TCP acks ? Assume infinite traffic on top of TCP and no packet loss/bottlenecks elsewhere. *Hint: Exploit TCP header fields & ack clocking***

- **3. Why are “sequence numbers”, “checksums” and techniques like “aging” not required in RIP ? Why is “count-to-infinity” absent in OSPF? Why can transient loops form in OSPF?**

- **4. Default-free routing tables and policy routing are important problems of inter-domain routing protocols. Does BGP-4 address these problems ? If so, how ?**

- **5. Mention two ways in which the naming hierarchy different from the addressing hierarchy, given that names are of variable length and can be organized independent of addresses ? How is DNS (based on this hierarchy) scalable and yet efficient ?**

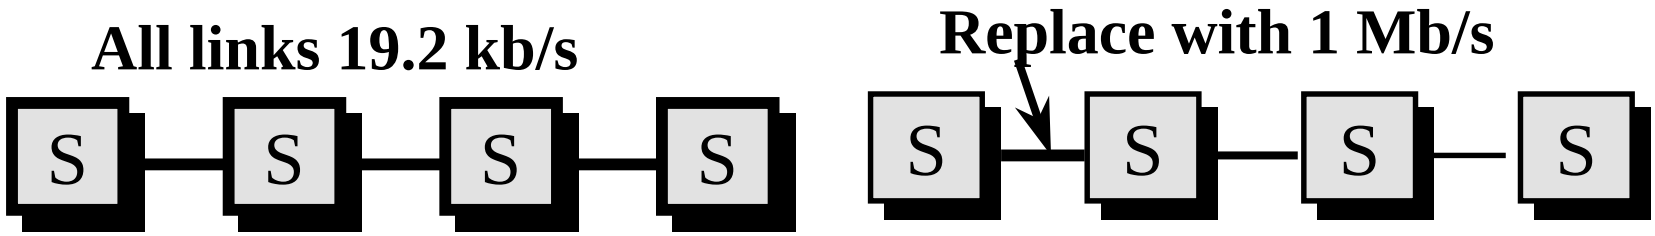

File Transfer time = ***5 mins***

File Transfer Time = ***7 hours***

- **1. Explain the problem in the figure to the right. If any other link were replaced by the 1 Mbps link, would you still encounter the congestion problem ? Why ? What does this tell you about placement (provisioning) of high speed links in a network ?**

- ***A) The problem of congestion arises because the high speed access link overloads the next (low speed) link. High overload leads to packet losses, timeouts, retransmissions, loss of retransmitted packets, timer backoffs etc slowing down information transfer substantially.***
- ***B) No. Because there is no overload situation assuming traffic flows from left to right.***
- ***C) Lesson: Access links should be low speed, and high speed links should first be deployed in the core.***

- **2. We would like to control the rates of TCP flows using a network component which can manipulate TCP acks. You are given the desired rate (R) and round trip time (RTT) for a TCP flow. How would you enforce this maximum rate and spread packet transmissions over an RTT by only manipulating TCP acks ? Assume infinite traffic on top of TCP and no packet loss/bottlenecks elsewhere.** ***Hint: Exploit TCP header fields & ack clocking***

- ***We need to translate the rate into a window value because TCP flow is controlled by a window variable = min (cwnd, rcvr_wnd)***
- ***Compute: W = RTT*R. Overwrite the receiver window field in the acks by this value. Recompute & rewrite TCP checksum.***
- ***Now, given this window value (written in acks) and the knowledge that TCP rate at a given window is determined by the ack rate, release W/MSS acks spaced uniformly over an RTT. Acks need to be buffered and delayed till this point. This will spread out packet transmissions uniformly over RTT.***
- ***FYI, Packeteer (packeteer.com) makes such a product.***

- **3. Why are “sequence numbers”, “checksums” and techniques like “aging” not required in RIP ? Why is “count-to-infinity” absent in OSPF? Why can transient loops form in OSPF?**
- *RIP uses a distance vector approach where it requires information about neighbors’ signposts (distance vectors), and the cost of adjacent links. Sequence numbers are required for duplicate detection, out-of-order detection etc. RIP information traverses just one link (which does not duplicate/reorder information) and therefore sequence numbers are unnecessary. Aging is used to deal with wrap around of sequence numbers -- not required when sequence numbers are not used. RIP runs over UDP and relies on the UDP checksum (OSPF runs over IP directly)*
- *The count-to-infinity problem occurs because some vector element increases unboundedly due to lack of complete path information. OSPF reliably gathers an entire map at every node and knows complete path information. Therefore, loops cannot be formed by count-to-infinity. However, due to transient delays in transmitting a larger amount of state information over unreliable channels, inconsistent information can result in different nodes leading to loops.*

- **4. Default-free routing tables and policy routing are important problems of inter-domain routing protocols. Does BGP-4 address these problems ? If so, how ?**

- ***BGP-4 uses a path-vector based distance vector approach. Every autonomous system is given an AS#, and reachability to (address prefixes in) AS's is represented in terms of a path-vector of AS#s.***
- ***Default-free tables are maintained as follows. Each AS is responsible for advertising reachability info for every address prefix contained in it. Every inter-domain router collects and maintains reachability to all such prefixes, and hence are default-free.***
- ***Policy involves constraining the set of ingress/egress routers for subgroups of transit traffic. This is accomplished through filtering of routing advertisements and/or use of "communities" attribute in BGP advertisements.***

- **5. Mention two ways in which the naming hierarchy different from the addressing hierarchy, given that names are of variable length and can be organized independent of addresses ? How is DNS (based on this hierarchy) scalable and yet efficient ?**

- ***The naming hierarchy follows an organizational (not routing/topological) hierarchy. Also there is no limit to the depth/breadth of this hierarchy (unlike addressing hierarchy which is limited by the finite length of the address).***
- ***DNS scales because DNS servers resolve only part of the name space and are organized in a hierarchy closely paralleling the upon the naming hierarchy. One potential problem with unbounded levels in the hierarchy is the number of hops which need to be traversed for resolving a name. This process is speeded up through caching, and use of non-authoritative replies.***