Internet Protocols: Quiz 2

- This quiz consists of true/false questions for 20 pts and four short answers/quantitative problems for 30 pts, a total of 50 points.
- In the True/False questions, the following grading policy will be used:
  - **Correct answer: +1 pt**
  - **Wrong or Blank/Unattempted answer: 0 pts {change in policy}**
- There will be ***no negative grading for the short answer or quantitative problems***. Partial credit may be awarded where appropriate.
- Open book policy
- ***Time: 60 min (one hour).*** **Strictly enforced.**
- This is the second quiz out of three quizzes. ***Best two out of three*** will be considered for final grades. Each of the two quizzes chosen will be weighted equally.
- NOTE: There is a small change in the Dijkstra's algorithm (given as handout for your reference - problem 4)

True or False? ***(20 points)***

Note: Correct Ans = +1; Wrong Answer or Did not attempt = 0

T F

❑ ❑ Since TCP may send two packets into the network for every ack, it does not aim to achieve the “conservation of packets” principle given in Van Jacobson’s paper.

❑ ❑ Load-balancing done by OSPF is not a good solution for congestion because it does not reduce the number of packets (demand) in the network.

❑ ❑ The current window size in TCP equals the congestion window (cwnd) variable.

❑ ❑ Use of the timestamp option in TCP would solve the retransmission ambiguity problem discovered by Karn and Partridge.

❑ ❑ If the SACK option is used, the TCP sender can advance its snd_una variable upon receiving a SACK and need not wait for an ACK.

❑ ❑ Nagle’s algorithm combined with delayed-acks solve the small packet (or tinygram) problem in TCP.

❑ ❑ When you see your idle ftp session reset by the server, it is because the server detects the underlying idle tcp connection using the TCP persist timer.

❑ ❑ T/TCP sacrifices reliability of transaction processing for speed - your bank transactions executing over T/TCP are unreliable.

❑ ❑ If I have a domain and want to add a new subdomain, I must register this subdomain with Network Solutions Inc., which handles Internet Name registry.

T F

❑ ❑ The window scaling option addresses the issue of filling up the large bandwidth-delay pipe.

❑ ❑ When an ICMP packet is dropped, an ICMP error message is sent to the router which generated it.

❑ ❑ The OSPF header contains a checksum because it runs directly over IP (RIP headers don't have checksum fields).

❑ ❑ The split-horizon fix solves the count-to-infinity problem in RIP.

❑ ❑ The sequence number and age fields are used in OSPF to optimize on the number of packets flooded into the network.

❑ ❑ VLSM support is why RIP is preferred over OSPF in small autonomous systems

❑ ❑ The "core" architecture with EGP was a scalable solution to the Internet routing problem.

❑ ❑ Policy routing cannot use the Bellman-Ford algorithm because multiple metrics are used and the least "cost" route need not be the desired route.

❑ ❑ Multiple CIDR routing entries may match a given destination network address.

❑ ❑ If the DNS server cannot resolve a recursive query, it will contact other servers and ultimately return an authoritative reply.

❑ ❑ If an AS sets a policy not to allow transit traffic through it, it is called a stub-AS.

- 1) (7 pts) Appendix A of Van Jacobson's paper talks about the problem with the RFC 793 RTT estimation algorithm. Why is the old algorithm bad and why is the use of standard deviation useful in the calculation ? Why is the mean deviation used instead of standard deviation ?

- 2) (7 pts) State briefly how Fast-retransmit-and-recovery (FRR) and Selective Acknowledgements (SACK) improve the basic TCP congestion control algorithm.

- 3) (7 pts) The command 'nslookup midway.cis.ohio-state.edu' generates a DNS query for the name 'midway.cis.ohio-state.edu'. Describe how the DNS system would resolve this query if I do this on our experimental machine, beethoven.

- 4) (9 pts) The LSP database at node 2 for a 3-node network (shown below) is as follows. Notation: (Link: Cost)

(1-2: 2), (1-3: 7), (2-1: 4), (2-3: 2), (3-1: 1), (3-2: 2). The sets PATH and TENTATIVE are initially empty. Trace the steps of Dijkstra algorithm (see handout) at node 2 to find the least cost routes from node 2 to nodes 1 and 3.

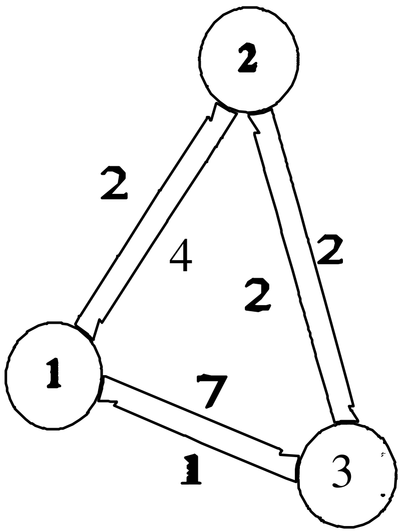

Dijkstra's algorithm(pblm:4)

- Assume that the following databases are present:
  - LSP database, set of (Link:Cost) pairs for every Link.
  - PATH = (ID, path cost, forwarding direction) triples. Set of nodes for which the best path's cost & forwarding direction have been found
  - TENTATIVE: (ID, path cost, forwarding direction) of candidate triples from which least-cost triples are moved to PATH
- Algorithm:
  - Start with "self" as the root of a tree: I.e put (myID, 0, 0) in PATH.
  - If N is just put into PATH, look at N's links (ie its LSP). For each link (to say neighbor) M, add cost from root-

Dijkstra's algorithm (contd)

to-N to the cost of link from N to M. If M not already in PATH or TENTATIVE with better path cost, ***add*** (M, new path cost, M-N link) to TENTATIVE. ***If triple for node M, (M, old path cost, old link), is already present in TENTATIVE, replace it with the new triple (M, new path cost, M-N link).***

- If TENTATIVE empty, terminate. Else, take the min-cost triple from TENTATIVE and move to PATH. ***Again replace triple in PATH if necessary as in step 2.*** Go to step 2.

True or False? ***(20 points)***

Note: Correct Ans = +1; Wrong Answer or Did not attempt = 0

T F

❑ √ Since TCP may send two packets into the network for every ack, it does not aim to achieve the "conservation of packets" principle given in Van Jacobson's paper.

√ ❑ Load-balancing done by OSPF is not a good solution for congestion because it does not reduce the number of packets (demand) in the network.

❑ √ The current window size in TCP equals the congestion window (cwnd) variable.

√ ❑ Use of the timestamp option in TCP would solve the retransmission ambiguity problem discovered by Karn and Partridge.

❑ √ If the SACK option is used, the TCP sender can advance its snd_una variable upon receiving a SACK and need not wait for an ACK.

❑ √ Nagle's algorithm combined with delayed-acks solve the small packet (or tinygram) problem in TCP.

❑ √ When you see your idle ftp session reset by the server, it is because the server detects the underlying idle tcp connection using the TCP persist timer.

❑ √ T/TCP sacrifices reliability of transaction processing for speed - your bank transactions executing over T/TCP are unreliable.

❑ √ If I have a domain and want to add a new subdomain, I must register this subdomain with Network Solutions Inc., which handles Internet Name registry.

T F

√ ❑ The window scaling option addresses the issue of filling up the large bandwidth-delay pipe.

❑ √ When an ICMP packet is dropped, an ICMP error message is sent to the router which generated it.

√ ❑ The OSPF header contains a checksum because it runs directly over IP (RIP headers don't have checksum fields).

❑ √ The split-horizon fix solves the count-to-infinity problem in RIP.

√ ❑ The sequence number and age fields are used in OSPF to optimize on the number of packets flooded into the network.

❑ √ VLSM support is why RIP is preferred over OSPF in small autonomous systems

❑ √ The "core" architecture with EGP was a scalable solution to the Internet routing problem.

√ ❑ Policy routing cannot use the Bellman-Ford algorithm because multiple metrics are used and the least "cost" route need not be the desired route.

√ ❑ Multiple CIDR routing entries may match a given destination network address.

❑ √ If the DNS server cannot resolve a recursive query, it will contact other servers and ultimately return an authoritative reply.

√ ❑ If an AS sets a policy not to allow transit traffic through it, it is called a stub-AS.

- 1) (7 pts) Appendix A of Van Jacobson's paper talks about the problem with the RFC 793 RTT estimation algorithm. Why is the old algorithm bad and why is the use of standard deviation useful in the calculation ? Why is the mean deviation used instead of standard deviation ?
- Soln:
  - The old algorithm calculates the mean RTT and then sets RTO as β*RTO, where $\beta = 2$. Jacobson showed that this can tolerate variations of only upto 30%. In reality the RTT variances are far greater. Also the averaging technique is constrained to use a small weight, g.
  - Standard deviation is a well known statistical measure of average variation, which can be estimated along with the mean RTT. Once this is done, RTO = mean RTT + h*STD_DEV, would tolerate variations of a factor of h.
  - The mean deviation is shown to be a conservative estimate of standard deviation and is easy to calculate mathematically (only integer operations needed).

- 2) (5 pts) State briefly how Fast-retransmit-and-recovery (FRR) and Selective Acknowledgements (SACK) improve the basic TCP congestion control algorithm.
- Soln:
  - The first problem in TCP is that packet loss was detected by a timeout which could be very long (implementations used a timer granularity of 500 ms). Use of duplicate acks speeds up detection - this is the "fast retransmit" part of FRR. The second problem is the response to congestion - the additive increase/multiplicative decrease rule says that the window needs be cut down by a factor, followed by additive increase. However, TCP sets the window to one and performs slow start until it reaches SSTHRESH (a go-back-N retransmission strategy). This delay can be avoided - which is the fast recovery strategy.
  - The TCP with FRR is still vulnerable to timeout and "beat-down" of SSTHRESH when a burst loss occurs. SACK indicates which blocks have been received, allowing the sender to optimize on what to retransmit.

- 3) (5 pts) The command ‘nslookup midway.cis.ohio-state.edu’ generates a DNS query for the name ‘midway.cis.ohio-state.edu’. Describe how the DNS system would resolve this query if I do this on our experimental machine, beethoven.
- Soln:
  - The DNS resolver on beethoven would send a DNS query to the DNS server. If the DNS server has the IP address corresponding to this name (unlikely) it will resolve the name and send back a non-authoritative reply. Else, (assuming that the query is a recursive query) it will contact other servers and return a non-authoritative reply, and the name and addresses of the servers which the client could contact to confirm the reply.

- 4) (10 pts) The LSP database at node 2 for a 3-node network (shown below) is as follows. Notation: (Link: Cost)

(1-2: 2), (1-3: 7), (2-1: 4), (2-3: 2), (3-1: 1), (3-2: 2). The sets PATH and TENTATIVE are initially empty. Trace the steps of Dijkstra algorithm at node 2 to find the least cost routes from node 2 to nodes 1 and 3.

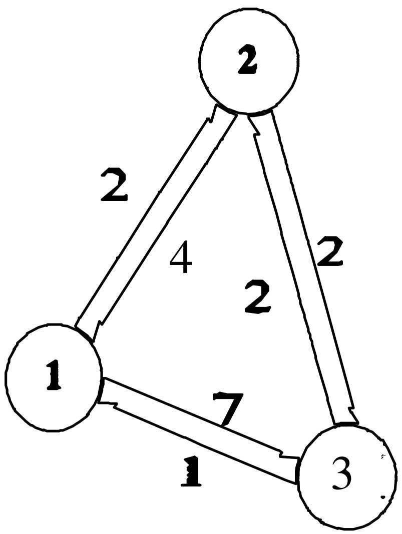

1. Add (2, 0, 0) to PATH. Look at neighbors of 2 I.e. (2-1: 4), (2-3: 2).

2. TENTATIVE = { (3, 2, 2-3), (1, 4, 2-1)}. Choose the least cost triple (3,2,2-3) and move to PATH, I.e., PATH = {(2, 0, 0), (3, 2, 2-3)}. Look at neighbors of 3, I.e. (3-1: 1), (3-2: 2).

3. (1, 2+1 = 3, 2-3) is better than (1, 4, 2-1) in TENTATIVE. Therefore TENTATIVE = {(1,3,2-3)} (replace (1,4,2-1) with (1,3,2-3) …). Since TENTATIVE contains only one triple, move it to PATH. PATH = {(2, 0, 0), (3, 2, 2-3), (1,3,2-3)} Examining neighbors of 1, we don’t add anything new to TENTATIVE => terminate.