

ECSE-6600: Internet Protocols

Exam 2

Time: **75 min (strictly enforced)**

Points: **50**

YOUR NAME:

*Be brief, but **DO NOT** omit necessary detail*

{Note: Simply copying text directly from the slides or notes will not earn (partial) credit. Brief, clear and consistent explanation will.}

I. Below, you are given a true or false statement and asked a follow up question.

1. [5 pts] **True statement**: TCP uses a three-way handshake for connection setup.

(5 pt) Explain why TCP does not use a two-way handshake. What extra functionality is gained in a three-way handshake that is not possible with a two-way handshake?

3-way handshake makes duplex connections possible. We need 2 sets of sequence numbers for two-way communication.

2. (5 pts) **True statement**: The IP address, Area ID, AS number are all identifiers.

(5 pts) Explain why we need new identifiers like area ID and ASN? Why not encode them into the IP address, similar to the way a network ID and interface ID were encoded ?

There can be a lot of areas with common IP prefixes. So identifying areas with IP addresses will be hard. Also a single area might encompass many IP prefixes.

Routing among AS's becomes possible.

3. (5 pts) **True statement**: OSPF uses a sequence number field and an Age Field
(5 pts) Explain why we need both these fields ? Why not solve the problems which they address with a single field?

Age prevents old LSAs from looping in the network forever or to avoid stale information from affecting the database. This is similar to TTL field in IP packets. We need sequence number field to consider only the latest LSAs. If we have just a sequence number field, we will not be able to remove looping LSA packets. If we have only the age field, we will not detect duplicates.

4. (5 pts) ***True statement:*** Hello exchange and database synchronization is done using different methods when OSPF is mapped to a broadcast LAN, compared to how it is done on a point-to-point link

(5 pts) Explain what are the differences in how Hellos are exchanged and how database synchronization is done in broadcast LANs, when compared to how it is done on a point-to-point link.

In a broadcast LAN:

A Hello is sent to the multicast address “AllSPFRouters” and this prevents re-advertisement, since all routers on the broadcast link receive the Hello. Instead of sending N-1 (assuming there are N hosts in the network) Hellos we send only 1 Hello.

LSAs are sent to the multicast address “AllSPFRouters” and the LSA ACKs are sent only to “AllDRRouters” address. Hence this avoids sending separate copies to DR and BDR.

5. (5 pts) **True statement:** EGP and I-BGP can result in looping packets if a full mesh or tree structure is not imposed, whereas E-BGP does not have this problem

(5 pts) Explain why EGP and I-BGP have this looping problem, and why E-BGP does not have this problem.

Header info in EGP and I-BGP does not give distance or a list of path elements that can prevent a loop. Therefore a logical topology that is loop-free has to be enforced.

6. (5 pts) **True statement**: Recursive lookup is performed by BGP before accepting any route.

(5 pts) Explain why recursive lookup is necessary in BGP (it was not needed in OSPF!)?

BGP nexthop is not the same as IP-nexthop and doesn't guarantee the existence of an IP-nexthop. Recursive lookup resolves BGP-nexthop to give the IP-nexthop.

II. [10 pts] TCP Congestion Control

(4 pts) TCP congestion control has four key pieces: self-clocking, slow start phase, congestion avoidance phase and fast retransmit/recovery. Explain what is the role of each piece, and how they are complementary to each other in terms of functionality.

(3 pts) What are the key differences between TCP Reno and SACK from both a retransmission perspective and congestion control perspective?

(3 pts) How does the use of active queue management (AQM) schemes like random early detection/drop (RED) at routers help TCP congestion control over-and-above the congestion control functions at end-systems?

Self-clocking: Takes care of packet conservation principle, that is the network has as many unacked packets as the window provides.

Slow-start: Probes the available bandwidth to the connection quickly using an exponential window increase.

Congestion Avoidance: In this phase, congestion has been detected and the window increase is much slower to avoid further losses.

Fast retransmit/recovery: To facilitate quick detection of congestion, TCP reacts to triple duplicate acks by cutting its window to ssthresh. It retransmits the packet for which the ACK was not received – this is fast retransmit; by fast recovery, we mean the window is not allowed to fall to 1, TCP is allowed to send a packet if the window allows it.

Reno vs SACK: In Reno we have cumulative ACKs and so a packet loss causes retransmit starting from the unacked packet. With SACK all the received bytes are acked. So the retransmission is only for those bytes that were lost. From a congestion control perspective SACK avoids more timeouts.

AQM: helps TCP reduce losses by a) Conveying early congestion signals, b) Distributing losses randomly and hence preventing burst losses and consequent timeouts.

IV. [10 pts] BGP Policy Routing and Load Balancing:

- a) (5 pts) BGP allows for policy routing by filtering routes at the inbound and outbound points. Discuss the advantages and limitations of this approach, compared to a link-state approach where each AS advertises the policies it has for every inter-domain link, and each node uses this “map” to determine an end-to-end route.
- b) (5 pts) BGP provides multiple hooks for load-balancing: MED, LOCAL-PREF, punching holes in prefixes/subverting CIDR, AS-path padding and communities. Briefly state some goals for load-balancing and analyze how effective each of these hooks are in meeting those goals.

a)Advantages:

Filtering is a local strategy and changes can be effected quickly. It has no control overhead since we need not advertise the policy.

Limitations:

Lack of global information for computing routes could mean sub-optimal routes being chosen.

Policy can be anonymous (not everyone knows my policy).

b)Goals for load balancing:

To equally distribute inbound/outbound traffic among multiple paths available. To avoid route-flapping while still adjusting routing to effect load balancing.

MED can be used to specify preferences for incoming traffic. But they can export instability.

Local-pref can help choose between multiple exits from an AS, but this might override an external AS’s preferences.

Subverting CIDR can help receive traffic on different interfaces, but this increases routing table sizes remotely and changes cannot be effected quickly.

Padding can regulate inbound traffic by advertising artificially larger AS-Paths.

Communities allow arbitrary local inbound/outbound policies given that other providers co-operate.

