SOLUTIONS TO CS 536 FINAL, FALL 2021 (PARK)

P1(a) 20 pts
3/3/3/3/3/5

Slow compared to sending data all at once.
3 pts

Exponential increase is considered fast, not slow.
3 pts

Probe quickly (exponential increase) to roughly estimate available bandwidth
which is detected by triggering timeout.
3 pts

Transfer of small files will likely be completed while in Slow Start phase.
3 pts

Congestion Avoidance uses linear increase to approach the maximum throughput
achievable after Slow Start.
3 pts

Congestion means that if offered load is excessive throughput starts to decrease.
When sending rate is reduced after detecting congestion, load may continue
to increase since throughput has decreased as well. Exponential decrease makes
it likely that load will decrease. Linear decrease may not.
5 pts

P1(b) 20 pts
4/2/2/2/4/2/4

Company conveys to its registrar (e.g., GoDaddy.com): primary (ns1.newkiddglobal.com)
and secondary (ns2.newkiddglobal.com) DNS servers and their IP addresses.
4 pts

Company adds type A record for IP address of web server www.newkiddglobal.com
[possibly add newkiddglobal.com as shortcut for www.newkiddglobal.com]:
<www.newkiddglobal.com, a.b.c.d, A>. Must specify TTL (time-to-live in seconds)
which is omitted for brevity here and below. a.b.c.d is meant to be an IPv4
address. For simplicity of illustration we assume IPv4 addresses.
2 pts

Company adds type MX record for mail server mymail.newkiddglobal.com ("mymail" is
just an example name): <newkiddglobal.com, mymail.newkiddglobal.com, MX>.
2 pts

Company adds type A record for IP address of mail server mymail.newkiddglobal.com:
<mymail.newkiddglobal.com, m.n.o.p, A> where m.n.o.p is the IP address of the
mail server.
2 pts

Method 1: create an alias using type CNAME:
<www.newkiddglobal.com, s1.newkiddglobal.com, CNAME> and type A record for the
new web server: <s1.newkiddglobal.com, a.b.c.d, A>.
4 pts

Method 2: modify type A record of www.newkiddglobal.com at primary and secondary
name servers to: <www.newkiddglobal.com, e.f.g.h, A> where e.f.g.h is the IP
address of s1.newkiddglobal.com. [Optionally add type A record specifying IPv4
address of s1.newkiddglobal.com as e.f.g.h.]
2 pts

Delay in global DNS resolution to the new server IP address depends on the TTL
values of www.newkiddglobal.com's type A record. Cached values that are stale will
be used until TTL expires.
4 pts

P2(a) 15 pts
5/2/3/3/2

OSPF: Dijkstra shortest-path algorithm, delay as performance metric, uses IP

RIP: Bellman-Ford shortest-path algorithm, hop count as performance metric, uses UDP
IS-IS: Dijkstra shortest-path algorithm, delay as performance metric, uses Ethernet
// also ok to state broadcast for OSPF/IS-IS instead of Dijktra, and nearest-neighbor communication instead of Bellman-Ford
5 pts

ISPs with large intra-domain networks use OSPF/IS-IS because RIP using Bellman-Ford is slow to converger. Purdue has a small intra-domain network which is suited to running RIP.
2 pts

Z's BGP border router should communicate before the start of daytime to X's border router that Z's IP prefixes should be routed through X. At the end of daytime, Z's border router communicates to X's border router not to route Z's prefixes through X.
3 pts

Z's BGP border router should communicate before the start of nighttime to Y's border router that Z's IP prefixes should be routed through Y. At the end of nighttime, Z's border router communicates to Y's border router not to route Z's prefixes through Y.
3 pts

Z's BGP router sends an AS-PATH update message at the beginning of daytime to X's border router and at the beginning of nighttime to Y's border router. Z's BGP router sends a withdraw message at the end of daytime to X's border router and at the end of nighttime to Y's border router. The AS-PATH update message should be sent before the withdraw message during switch-over so that Z remains reachable on the global Internet.
2 pts

P2(b) 15 pts
3/3/3/3/3

Customers at coffee shops J, K, and L get about the same throughput.
3 pts

Since a station at one of the three hot spots can sense a station at any of the three hot spots, carrier sense (CS) is the same across the three coffee shops and resultant throughput approximately the same.
3 pts

Customer at coffee shops J and L get about the same throughput whereas customers at K get less throughput.
3 pts

A station at J can carrier sense stations at J and K, but not L. Similarly, a station at L can carrier sense stations at L and K, but not J. In contrast, a station at K can carrier sense stations at all three coffee shops. Since CS prevents a station from transmitting (CA of CSMA), stations in the middle coffee shop K will get less opportunity to transmit resulting in less throughput compared to J and L.
3 pts

A simple solution is for the three coffee shops to use three different, nonoverlapping channels.
3 pts

P3(a) 15 pts
4/4/4/3

Enabling DiffServ on its routers alone will not provide meaningful differentiated services to its customers since bottlenecks may occur at routers not belonging to the ISP where DiffServ remains disabled.
4 pts

Priority scheduling suffers under starvation. WFQ guarantees a minimum service level to each service class thus preventing starvation.
4 pts

No. Even if an IPv4 packet whose TOS field is marked platinum enters a router with DiffServ enabled where platinum packets get 60% of a link's bandwidth and silver packets get only

10%, if volume (or rate) of traffic marked as platinum is significantly higher than those marked as gold, platinum packets may receive inferior service.
4 pts

If DiffServ were implemented using priority scheduling then packets marked platinum would also receive better service than packets marked silver, irrespective of the volume (or rate) of packets belonging to the two classes.
3 pts

P3(b) 15 pts
11/4

Assuming buffer occupancy is at Q* when a packet loss occurs. From this moment, Q* times 0.1 sec elapses before the missing packet must be played on the audio device. Since timeout is 2 RTT and one-way delay RTT / 2, from this moment 2 RTT (= 1.5 RTT to trigger timeout plus 0.5 RTT for one-way delay) will elapse before the retransmitted packet reaches the client. Thus 0.1 Q* >= 2 RTT, i.e., Q* >= 20 RTT, suffices for retransmitted packets to arrive on time for playback on the audio device.
11 pts

Priority queue where packets are sorted by sequence numbers.
4 pts

Bonus 10 pts
2/2/2/1/1/2

Persistent TCP connection to serve multiple HTTP requests.
2 pts

Opening/closing separate TCP connections for each HTTP request incurs significant overhead due to the overhead associated with TCP connection set-up (and tear down).
2 pts

In HTTP/1.1 where multiple GET requests are sent sequentially over the same TCP connection, the response from the web server must arrive in the same order. This implies that a GET request requiring significant time delays short GET requests that are behind it.
2 pts

Several threads (e.g., 6) each using its own TCP connection can help partly mitigate HOL.
1 pts

However, if the number of GET requests are many with several incurring significant delay, even with multiple threads a thread can suffer under HOL.
1 pts

The main technique used by HTTP/2 is framing where HTTP responses are divided into smaller units (i.e., frames) that can be interleaved when sent over the same TCP connection and reassembled correctly at the client. Hence HOL is prevented.
2 pts