SOLUTIONS TO CS 536 FINAL, FALL 2016 (PARK)

P1(a) 10 pts

Intra-domain: routing decisions within an AS's intranet.
Inter-domain: routing decisions between autonomous systems.
2 pts

Main difference: intra-domain uses shortest path routing,
inter-domain follows policy routing.
2 pts

Suppose host A at AS X sends a packet destined to host B on
AS Y. Intra-domain determines path within AS X to a border
router R. At R, policy routing determines which AS to traverse.
Intra- and inter-domain routing repeats until packet reaches
a border router at AS Y from which intra-domain routing delivers
the packet to host B.
2 pts

LANs can be comprised of multiple switches. With IP turned off
for efficiency, LAN routing carries out the responsibility of
forwarding a LAN packet to its destination.
2 pts

There must not be a loop which is achieved through a protocol
that builds a spanning tree over the LAN switches.
2 pts

P1(b) 10 pts

GET /index.html http/1.1\r\n
host: www.purdue.edu\r\n
\r\n
[the spaces above represent spaces in ASCII]
2 pts

version 1.0 terminates TCP connection after every client request
such as GET.
1 pts

A header line "cookie: name=xyz" is included.
1 pts

A service provider can use the cookie value in subsequent GET requests
to track clients and provide individualized service.
2 pts

Suppose a HTTP client has cached a previous response from a server along
with the time stamp returned. By transmitting the saved time stamp as
the value of the header line containing if-modified-since, if the cached
copy is current (i.e., has not been changed at the server) then the
server can reply with status code 304 indicating no change which reduces
response time and traffic overhead.
2 pts

By transmitting the ETag value returned by a server in the next GET request
with header line containing if-none-match, the server can reply with status
code 304 in case the requested URL has not changed.
2 pts

P1(c) 10 pts

Indoors the effects of spatial diversity stemming from constructive/destructive
electromagnetic (EM) wave interference is pronounced. That is, signal strength
at one location may be strong if multipath reflections constructively interfere,
weak if they do so destructively.
3 pts

In general, no. Although signal strength deteriorates (i.e., attenuates) with
distance, spatial diversity can render a location x that is farther in distance
from a location y to a base station/access point stronger with respect to
signal strength/reception.
3 pts

Move location.
Change frequency/channel.
4 pts

P1(d) 10 pts

The main cause is heavy-tailed file sizes (i.e., mice and elephants).

2 pts

Compressed video operates by utilizing similarities across frames within a scene (inter-frame compression). As such, a reference frame within a scene may not be significantly compressed, however, other frames that differ from the reference frame in minor ways can be significantly compressed. This leads to significant variation in compressed video frame sizes.
2 pts

Compressed video is much burstier than compressed voice.
Compressed voice has a much smaller data rate (i.e., bandwidth requirement) than compressed video.
2 pts

IETF int-serv which is reservation-based.
IETF diff-serv which uses labels inscribed in the TOS field to provide differentiated services (e.g., platinum, golf, silver, bronze) to packets at a router.
2 pts

The main challenge is the need for a global infrastructure that enforces orderly access. Otherwise, users may tag their packets as high priority (e.g., platinum) which renders platinum meaningless if everyone does so or reserve bandwidth at will. Access control with pricing or other means to regulate access is required.
2 pts

P2(a) 10 pts

Congestion Avoidance implements linear increase/exponential decrease (i.e., method B) that conservatively aims to utilize available bandwidth.
Slow Start increases the congestion window exponentially fast in order to rapidly utilize network bandwidth and gauge where the congestion point may lie.
2 pts

A misnomer in the sense that exponential increase cannot be considered slow. However, compared to sending all at once, it may be viewed as slower.
2 pts

Unless a network is heavily congested, most files -- due to their small size -- complete transmission during the initial Slow Start phase. That is, Congestion Avoidance may not come into play.
For large files, Congestion Avoidance is dominant, intersperesed with Slow Start when timeouts occur.
4 pts

When duplicate ACKs arrive at the sender, it means that packets (e.g., belonging to a packet train) have made it to the receiver to trigger ACKs, however, one or more losses have occurred that set the sequence number of the ACKs to the first missing packet. Assuming the losses are spurious (and not due to network congestion), Fast Retransmit does not wait timeout of the missing packet but retransmits the missing packet when three duplicate ACKs are received.
2 pts

P2(b) 10 pts

Using OFDMA for multi-user communication would entail allocating/reserving carrier frequencies to individual users before data can be transmitted which incurs significant overhead. In WLANs that are not overloaded, CSMA (i.e., DCF) promotes bandwidth sharing without OFDMA's reservation overhead while achieving adequate performance.
4 pts

Also, in the case where there are many users associated with an access point (i.e., the number of users exceeds the number of available frequencies) but each user is using bandwidth intermittently, using CSMA instead of OFDMA can accommodate more users.
2 pts

In WLANs, especially in indoor environments, spatial diversity renders some locations better than others with respect to signal strength due to constructive/destructive interference. A location that receives poor signal reception at carrier frequency f may receive good signal strength at another carrier frequency f'. By spreading the data bits of a single user's packet across multiple carrier frequencies in conjunction with FEC, the bits being corrupted at f may be recovered by the successfully received bits carried on other carrier frequencies due to FEC.
4 pts

P2(c) 10 pts

QoS routing means routing with two or more routing metrics such as bandwidth and delay all of which must be satisfied by a selected path.
2 pts

QoS routing is a combinatorial optimization problem that is known to be NP-hard. In

practice, it is not considered a significant issue because currently there is not much need for QoS routing and if there is a need, heuristics/approximation algorithms exist that give suboptimal but adequate solutions.
2 pts

OSPF runs over IP, IS-IS runs directly over the underlying link layer protocol (e.g., Ethernet).
1 pts

RIP implements Bellman-Ford whereas OSPF implements Dijkstra (for shortest path).
1 pts

Purdue used RIP within its IP intranet for the longest time.
1 pts

Bellman-Ford which implements a distributed algorithm for shortest-path is known to be susceptible to convergence issues which increases with network size and tier-1 providers operate large intranets.
1 pts

Purdue is a stub autonomous system which means that Purdue does not provide transit service to packets originating from other autonomous systems.
1 pts

Purdue is known by IP address blocks advertised in CIDR format. They include: 128.10.0.0/16, 128.210.0.0/16, 204.52.32.0/20
1 pts

P3(a) 15 pts

In a distributed system with two parties (e.g., hosts), the aim of 2-party consensus is to have a protocol that after finite time assures that both parties come to agreement.
2 pts

The intuition is that in a 3-round handshake, the last ACK of the handshake may go missing and there is no way for the sender who sent the last ACK to know the status unless an ACK for the last ACK is received. This holds for any k-round handshake which indicates that agreement may not be reached in finite time.
2 pts

In TCP connection termination, both parties must come to agreement that their connection has terminated so that they can free up TCP related data structures allocated inside an operating system.
2 pts

TCP implements a heuristic/hack by making the party that sent the last ACK set a timer (e.g., two minutes) during which the connection remains live. In case, the last ACK was not received, this gives the other party an opportunity to request retransmission of the missing ACK.
2 pts

Two hosts A and B have established a TCP connection and perform data transfer. B sends a FIN packet to A indicating that it is done sending (TCP is full-duplex). A sends an ACK that acknowledges the FIN. If A is done sending data as well, it will piggyback the ACK with its own FIN packet. Otherwise, a FIN packet is transmitted when A is done sending data. B acks the FIN and sets a timer during which the connection remains live.
2 pts

If the last ACK from B to A is not received, A retransmits its FIN multiple times because the FIN packets get lost, by the time a retransmitted FIN reaches B, B's timer may have expired at which point B's operating system will have deleted all the data structures associated with the connection. Hence when the retransmitted FIN arrives at B, B will not be able to ack which results in A not being able to terminate the connection.
2 pts

Application programmers experience this issue when they close a TCP connection using a port P and try to reuse it immediately/soon thereafter but the operating system returns an error specifying that the freed port is not available.
1 pts

TCP connections are usually set up to transfer data. Although TCP's SYN packets cannot guarantee consensus that a connection has been established, the subsequent data packets and their ACKs perform this function since data packets are retransmitted until they are acknowledged.
2 pts

P3(b) 15 pts

Method A was not able to regulate the sending rate to achieve bounded oscillation on its own. That is, it relied on the boundary conditions that sending rate cannot be negative and the number of in-flight packets cannot be negative to achieve oscillation around the target Q*.
2 pts

Method B by implementing exponential decrease (linear increase remains the same) is able to

achieve bounded oscillation on its own without the help of the boundary conditions.
2 pts

Instead of converging to the target Q*, Method B kept oscillating around it.
2 pts

In Method C, the gap/distance from where the state (i.e., Q(t)) is and what the target Q* is is incorporated in adjusting the sending rate.
2 pts

Method D added a damping term to Method C (that also oscillated) which made it converge to the target.
2 pts

Adopting Method D for TCP is not a simple porting matter for two reasons: first, in TCP data traffic (as opposed to multimedia streaming traffic), in general, there is no well-defined target Q*. Instead the best sending rate must be inferred by observing at what rate a TCP connection experiences congestion. Congestion Avoidance with the help of Slow Start is one way of doing so.
4 pts

Two, reaching the best sending rate and remaining there is not simple since network conditions can change and the optimal operating point is not easy to maintain.
1 pts

Bonus 10 pts

The Internet's end-to-end paradigm captures an approach to designing the network and its services that reduces the role played by routers/switches at the expense of increased complexity incurred by end systems that carry out functions that routers/switches do not perform.
4 pts

ARQ/sliding window is implemented by the ultimate sender/receiver without the help of intermediate IP routers.
2 pts

Secure communication such as confidentiality (encryption/decryption) are performed by end systems without the help of IP routers.
2 pts

In IPv4 fragmentation, routers may fragment IP packets if they are too large to fit in the payload of a link layer frame. However, they are not allowed to reassemble IP fragments. A router fragmenting a too large IP packet is a necessity. Reassemling fragments, even though in some cases feasible, is not a necessity and incurs overhead. Thus reassembly is delegated to the final receiver.
2 pts