

NAME:

Login name:

Computer Science 461
Final Exam
May 22, 2012
1:30-3:30pm

This test has seven (7) questions, each worth ten points. Put your name on *every page*, and write out and sign the Honor Code pledge before turning in the test. You should spend roughly fifteen minutes per question.

"I pledge my honor that I have not violated the Honor Code during this examination."

<i>Question (points)</i>	<i>Score</i>
<i>1 (10 pts)</i>	
<i>2 (10 pts)</i>	
<i>3 (10 pts)</i>	
<i>4 (10 pts)</i>	
<i>5 (10 pts)</i>	
<i>6 (10 pts)</i>	
<i>7 (10 pts)</i>	
<i>Total:</i>	

QUESTION 1: *Peer Pressure* (10 points)

1(a) The Skype peer-to-peer IP telephony application sometimes has a user machine serve as a “super node” to support communication between two other users. When are super nodes needed, and why? (3 points)

Traversing NATs and firewalls, particularly when both peers lie behind a NAT/firewall that blocks unsolicited traffic from external hosts.

Super nodes are also useful for discovering the remote peer.

1(b) Microsoft recently bought Skype. Earlier this month, Microsoft announced that they will replace the classic super nodes with their own super nodes running in Microsoft data centers. Give one reason why Microsoft might have made this decision? (2 points)

To improve performance and reliability by ensuring super nodes have sufficient resources and predictable operation.

1(c) To prevent free-riding, a BitTorrent peer rates other peers in terms of the download rate they offer, and then decides how much to offer in return. Suppose six users offer data at rates of 10, 4, 8, 3, 7, and 5 Mbps, respectively, and the peer has 15 Mbps of upload bandwidth. What upload rate does it offer to each of the six peers? (3 points)

The peer uploads with $1/N$ of its bandwidth to the top N peers, where N is 4 by default. So, the upload rates are $15/4$, 0, $15/4$, 0, $15/4$, $15/4$.

1(d) Under the same scenario as question 1(c), suppose peer #1 (currently offering a download rate of 10 Mbps) decides to be more strategic. What is the minimum rate that peer #1 could offer and still achieve the same upload rate it received in question 1(c)? (2 points)

Just slightly above the fifth peer, or just above 4 Mbps.

QUESTION 2: Insecurity (10 points)

2(a) Suppose one Autonomous System (AS) originates a bogus announcement for a prefix “owned” by another AS. Assume both the legitimate AS and the attacker are chosen at random among the ASes in the Internet. On average, what fraction of the ASes selects a bogus route to the attacker, rather than a legitimate one? Assume ASes do not defensively filter BGP messages, and that the owner and the attacker announce a prefix with the same mask length. (2 points)

Half, since the attacker and the victim are (on average) on equal footing to attract the traffic.

2(b) In Secure BGP, a BGP speaker v sends neighbor $a1$ an announcement for a route to destination prefix d . Speaker v uses its private key to sign a message that includes $a1$ (e.g., v signs the message “ $a1: (v, d)$ ”). Why does v include $a1$ in the signed message, instead of simply sending a signed version of “ v, d ” to $a1$? What attack does this prevent? (3 points)

If v only signs “ (v, d) ”, then $a1$ could give the signed route to other ASes to propagate as if v had sent the route to them directly.

2(c) Consider the protests in Egypt when the government shut down access to the Internet and the phone network. Suppose the protestors had built an “ad hoc” network to communicate in a peer fashion (say, using the Bluetooth interface on their phones). The protestors would want to sign and encrypt their messages. How could the users learn each other’s public keys, without access to a fixed network infrastructure or central certificate authority? (3 points)

The protestors come in direct physical contact, and can exchange keys in person (e.g., the “key signing parties” approach used in Web-of-trust).

2(d) In the TLS handshake protocol, the communicating hosts construct a shared symmetric key that they use to exchange data securely. Given two reasons why they don’t simply encrypt the data using asymmetric public-key cryptography. (2 points)

Symmetric key cryptography is much faster than asymmetric cryptography.

Clients often do not have a public-private key pair.

Repeatedly using the same keys increases the risk of compromise of the keys.

QUESTION 3: A Measured Response (10 points)

3a) Traceroute can sometimes return a path that does not even exist in the Internet topology. E.g., hop i may not be connected to hop $i+1$. Give one reason why this might happen. (2 points)

A routing change during the traceroute probe

Load balancers that spread traffic over multiple paths

3(b) Give an example where packet sampling may lead to an *increase* in the number of flows reported by Netflow. (2 points)

Sampling may cause an artificial increase in the inter-arrival time between packets of the same TCP/UDP connection, splitting a single connection into multiple flow records.

3(c) Suppose a Netflow log has three flows: (i) a 3-packet flow with a total of 300 bytes, (ii) a 5-packet flow with a total of 1000 bytes, and (iii) a 4-packet flow with a total of 1100 bytes. What is the average packet size across all traffic on the link? (2 points)

$$(300+1000+1100)/(3+5+4) = 2400/12 = \mathbf{200 \text{ bytes}}$$

3(d) On a *single* BGP session, why do BGP update messages for a single prefix often arrive close together in time? (2 points)

Path exploration during routing-protocol convergence.

3(e) Why do *different* BGP sessions (with different routers around the Internet) often have updates messages for the same prefixes at roughly the same time? (2 points)

The same routing change is observed from multiple vantage points, with many ASes switching to alternate routes for the same destination prefix at roughly the same time.

QUESTION 4: *Enterprising* (10 points)

4(a) Many enterprise networks only allow registered machines to access the local area network. Some of these networks configure the DHCP server to assign a computer (or MAC address) the same IP address every time it connects to the network. Give one reason why. (3 points)

Tracking traffic by IP address is much easier if an IP address corresponds to the same user at all time.

4(b) What is the *disadvantage* of associating each registered MAC address with a pre-determined IP address? (3 points)

*Running out of IP addresses when the number of **potential** users grows large, even if these users are not active at the same time.*

Less opportunity for subdividing the topology into separate IP prefixes for more scalable routing.

4(c) To send an IP packet to another host on the same local area network, the sending host first checks its local ARP cache to determine the MAC address associated with the destination IP address. On a “miss” in the ARP cache, the sending host sends an ARP request for the destination IP address and buffers the packet, awaiting a response. Suppose, instead, the host issues the ARP request and simultaneously *drops* the IP packet. Name one *advantage* of this approach? (2 points)

The sending host does not need to buffer the packet until receiving the ARP response. This reduces the memory requirements on the sender, and obviates the need to match the ARP response with some arbitrary packet sitting in a buffer.

4(d) What is one *disadvantage* of the approach in 4(c)? (2 points)

The dropped packets must be retransmitted at a higher layer (e.g., TCP). E.g., a dropped SYN packet, in particular, would have a long retransmission timeout.

QUESTION 5: *Going With the Flow* (10 points)

An OpenFlow switch has a prioritized list of packet-handling rules, each consisting of a pattern that matches on bits in the packet header and a set of actions to apply to the packets. This question explores applications of OpenFlow switches.

5(a) An OpenFlow switch can serve as an IP router. Suppose the “IP router” OpenFlow application has forwarding rules for three IP prefixes: (i) 12.1.2.0/24 forwards out link 1, (ii) 12.0.0.0/8 forwards out link 2, and (iii) 12.1.0.0/16 forwards out link 3. What *priority order* should these rules have in the OpenFlow switch, starting with the highest priority? (2 points)

The highest priority should go to the longest prefixes. That is, rule 1 should have priority over rule 3, which should have priority over rule 2.

5(b) Consider an OpenFlow switch where interface 1 connects to the local LAN and interface 2 connects to the rest of the Internet. Consider the design of a controller application that implements a simple repeater. What rule(s) should the controller install in the switch? For each OpenFlow rule, clearly indicate both the *pattern* and the *action*. (2 points)

pattern: input interface 1, action: output interface 2
pattern: input interface 2, action: output interface 1

5(c) Consider the design of an OpenFlow-based *stateful firewall*, where an external host cannot send traffic to an internal host unless the internal host contacts it first. To minimize the traffic unnecessarily sent to the controller, what are the initial rules in the switch, before any internal or external traffic arrives? For each rule, clearly indicate both the *pattern* and the *action*. (3 points)

pattern: input external-interface, action: drop
pattern: input internal-interface, action: send-to-controller

5(d) If an internal host with IP address 1.2.3.4 sends a packet to an external host with IP address 5.6.7.8, what new rule(s) should the controller install? For each rule, clearly indicate both the *pattern* and the *action*. (3 points)

pattern: src-ip 1.2.3.4, dst-ip 5.6.7.8, (optionally input internal-interface), action: forward external-interface
pattern: src-ip 5.6.7.8, dst-ip 1.2.3.4, (optionally input external-interface), action: forward internal-interface

QUESTION 6: *Wireless* (10 points)

6(a) Consider a wireless user connecting to the Internet via a wireless access point. Explain why collisions are *more* likely on a wireless network than a wired link, even if the two networks have the same link capacity and total demand. (2 points)

Hidden terminal problem

Collision avoidance rather than collision detection, causing collisions to send garbage on the network longer

6(b) Why does TCP congestion control perform poorly on wireless networks, compared to wired networks? (2 points)

TCP treats packet loss as a sign of congestion. However, in wireless networks, interference can cause packet loss. As such, the TCP sender sometimes erroneously drops its sending rate.

6(c) Suppose the access point runs a Web proxy. Explain why this may improve Web performance for wireless users, even if the cache miss rate is 100%. (3 points)

By terminating the TCP connection, the proxy separates the end-to-end path into two parts (i.e., the client-proxy and the proxy-server paths), and can retransmit lost packets separately on each part. For example, without the proxy, a packet from the server to the client could easily be lost on the wireless channel, forcing an end-to-end retransmission; with the proxy, the proxy could simply retransmit the data locally.

By “caching” a connection between the client and the proxy, the client avoids the repeated cost of TCP connection establishment.

6(d) Can a smart phone spread traffic to/from a single TCP connection over the WiFi and cellular interfaces at the same time? Why or why not? (3 points)

No, in today’s transport protocols, a transport connection is associated with a single IP address on each end-point, making it difficult to spread traffic over multiple interfaces connected to different administrative entities.

QUESTION 7: *Have Some Backbone* (10 points)

7(a) Give *two* reasons why an Internet Service Provider (ISP) might select a route with a longer AS-PATH over a route with a shorter AS-PATH. (2 points)

Avoids an undesirable AS (e.g., that performs censorship or wire-tapping)

Avoids an observed performance problem (e.g., high latency or low throughput)

7(b) A “tier 1” ISP is an AS with no upstream providers of its own. Why do the tier-1 ASes form a *clique* (i.e., a fully-connected topology), where each tier-1 ISP connects directly to every other tier-1 ISP? (2 points)

A tier-1 ISP has no upstream providers of its own. As such, it cannot reach the (single-homed) customers of other tier-1 ISPs without peering with them.

7(c) Suppose two neighboring ISPs have a “peer-peer” relationship with each other, and they peer in multiple locations. Why does one ISP require the other to announce prefixes at *all* of these locations, using BGP routes with the same AS-PATH length at *every* location? (2 points)

Otherwise, the ISP would not be able to perform “hot potato” routing to send traffic to the closest egress point connecting to the other ISP.

7(d) Suppose AS 1 has an AS-PATH of “1 2 3 4” to reach prefixes in AS 4, and ASes 3 and 4 have a peer-peer relationship. Suppose that an AS does not export BGP routes learned from one peer/provider to another peer/provider. Is AS 2 a customer, peer, or provider of AS 1? (2 points)

*AS 3 would only announce customer-learned routes to AS 4, implying that AS 2 is a customer of AS 3. Similarly, AS 2 would only announce customer-learned routes to AS 3, implying that AS 1 is a customer of AS 2. As such, AS 2 is a **provider** of AS 1.*

7(e) Give *two* reasons why interdomain routing uses *path*-vector routing instead of *distance*-vector routing. (2 points)

Faster convergence through loop detection, avoiding the “counting to infinity” problem

Flexible routing policies that depend on the hops in the AS-PATH.