Name:

SID:

Department of EECS - University of California at Berkeley
EECS122 - Introduction to Communication Networks - Spring 2005
Solution to the Final: 5/20/2005

There are 10 questions in total. Please write your SID on each page.

Question 1 (10 points)
The statements are true/false (T/F); mark the appropriate choice.

1. [X]T – [ ]F. Prefix-based addressing contributes to the scalability of the Internet.
2. [X]T – [ ]F. Packet-switching supports widely different connection types and link rates.
3. [ ]T – [X]F. The end-to-end principle is a set of rules for guaranteeing the quality of connections.
4. [ ]T – [X]F. The address resolution protocol is needed when different hosts have the same IP address.
5. [ ]T – [X]F. Generalized processor sharing is not used in routers because it is less efficient than weighted fair queuing.

Question 2 (10 points)

The following questions are multiple choice. Select the ONE answer that you think fits best.

1. Imagine ten packet flows, each controlled by a token bucket with size S bits and rate R bps. These flows go through one single router queue with bit rate C. The queue is initially empty. If $C > 10R$, then the maximum queuing delay cannot exceed
   (a) $S/(10R)$
   (b) 10S/C
   (c) S/R
   (d) None of the above.
2. Which one of the following is false?
   (a) Ethernet does not use a NAV (network allocation vector) in the header.
   (b) Wireless networks have higher propagation delay than wired ones.
   (c) Wired networks such as Ethernet do not have the hidden node problem.
   (d) Wired networks such as Ethernet do not have the exposed node problem.
3. Which of the following is a well-formed IP address?
   (a) inst.eecs.berkeley.edu
   (b) 00-E0-B8-49-52-67
   (c) 216.67.5.003
   (d) 264.112.34.9
4. To minimize the routing overhead in Ad-Hoc networks, following ideas are incorporated in the proposed routing protocols:
   (a) Perform routing on-demand basis.
   (b) Reduce the number of nodes involved in routing by delegating the networking routing function to only some of the neighborhood nodes.
   (c) First attempt to perform pro-active routing updates only along the optimal paths determined previously.
   (d) All of the above.
5. The RTS/CTS mechanism in **infrastructure** 802.11 Wireless LANs solves:
   (a) Exposed node problem.
   (b) Hidden node problem.
   (c) Both exposed and hidden node problems.
   (d) None of the above.

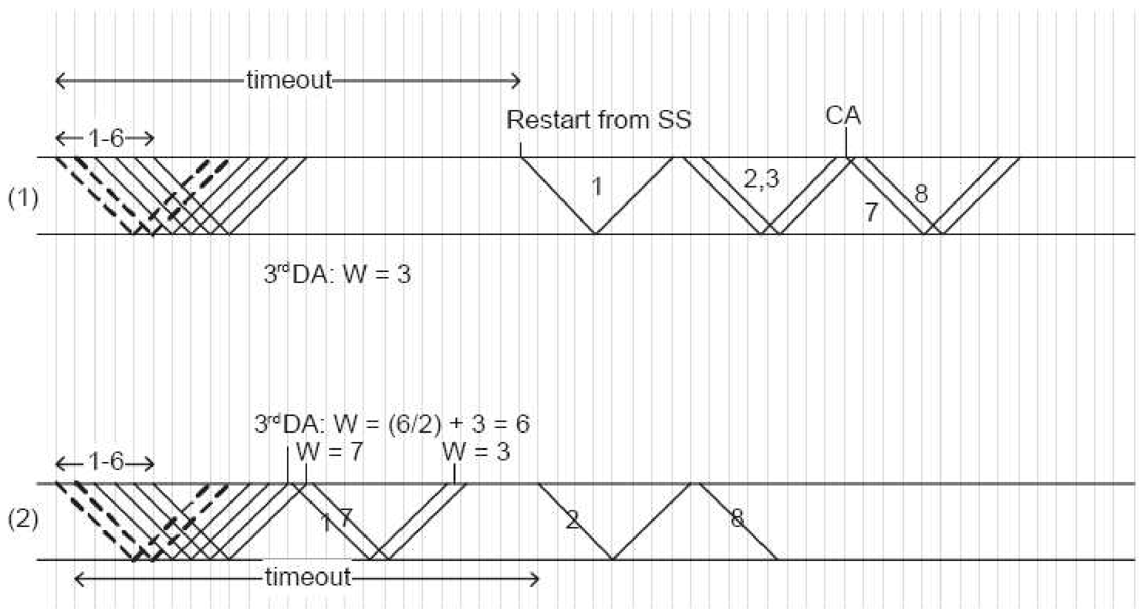

Figure 1: Diagrams for Question 3.

Question 3 (12 points)

Consider a TCP connection that, at time 0, is in the congestion avoidance phase with a window size equal to 6 MSS and then sends packets $\{1, 2, \ldots, 6\}$. (The connection has sent packets before that were properly acknowledged.) All the packets have a length equal to one MSS. The round-trip time is very large compared to one packet transmission time. Packets 1 and 2 are lost. No other packet or ACK is lost.

To be specific, let RTT = 8, transmission time = 1, timeout = 24.

(a) (6 points) Assume the sender does not implement fast retransmit nor fast recovery. After a timeout, it restarts in slow start. Draw a timing diagram that shows the transmission of the first eight packets by completing Figure 1 (1). Indicate the significant events, such as start of slow start (SS) or of congestion avoidance (CA) and the relevant congestion window sizes.

(b) (6 points) Repeat the previous problem assuming that the host uses fast retransmit and fast recovery. Complete Figure 1 (2). Indicate the significant events, such as start of slow start (SS) or of congestion avoidance (CA) and the relevant congestion window sizes.

Hint: When fast recovery begins, the window is set to CWND/2+3, and incremented by 1 with every subsequent duplicate ACK. When ACK for the retransmitted packet is received, the window is reset to half of the original CWND.

Solution

See Figure 1.

Question 4 (6 points)

Alice wants to send a document P to Bob. She wants Bob to be certain that the document has been sent by her. Explain how they can achieve that objective using public key cryptography.

Solution

One basic method is for Alice to send $C := D(P; Alice)$, i.e., the text obtained by applying the decryption algorithm specific to Alice. Bob can calculate $E(C; Alice)$ since the encryption rule for Alice is public. One has $E(C; Alice) = P$ because $E()$ and $D()$ are inverse functions of one another.

To be more secure, Alice could send $C := D(P * H(P); Alice)$ where $H(.)$ is an agreed upon hash function. This step reduces the possibility that Eve could calculate some $P' = E(C'; Alice)$ and pretend that she sent the text P' by transmitting $C' = D(P'; Alice)$. Indeed, it is very unlikely that Eve could find a C' with the property that $P' * H(P') = E(C'; Alice)$ for a meaningful P' that Eve would care about making Bob that Alice sent it.

Question 5 (15 points) Consider a commercial 10 Mbps Ethernet configuration with one hub (i.e., all end stations are in a single collision domain).

(a) (4 points) Find Ethernet efficiency for transporting 512 byte packets (including Ethernet overhead) assuming that the propagation delay between the communicating end stations is always 25.6 μs, and that there are many pairs of end stations trying to communicate.

Solution

Efficiency $= 1/(1 + 5a)$ where $a = 25.6 \cdot 10^{-6}/(8 \cdot 512 \cdot 10^{-7}) = 0.0625$

Efficiency $= 1/(1 + 5 \cdot 0.0625) = 0.76$

(b) (3 points) Now consider an application making use of TCP/IP/Ethernet for the same scenarios as in (a). Assume that the overheads of TCP, IP, and Ethernet are 20, 20, and 26 bytes, respectively. How much total bandwidth is available at the application layer in this network?

Solution

Total overhead = 66 bytes

Total application layer bandwidth available $= 10 \cdot 0.76 \cdot (512 - 66)/512$ Mbps $= 6.62$ Mbps

(c) (4 points) Recall that the maximum efficiency of Slotted Aloha is $1/e$. Find the threshold for the frame size (including Ethernet overhead) such that Ethernet is more efficient than Slotted Aloha if the fixed frame size is larger than this threshold. Explain why Ethernet becomes less efficient as the frame size becomes smaller.

Solution

We want $1/(1 + 5a) > 1/e$

$\Rightarrow 1 + 5a < e$ or $a < (e - 1)/5 = 0.34$

Let L bytes be the threshold frame size. Then, $25.6 \cdot 10^{-6}/(L \cdot 8 \cdot 10^{-7}) < 0.34$

$\Rightarrow L > 25.6/(0.34 \cdot 8 \cdot 10^{-1}) \approx 94$ bytes

(d) (4 points) Consider an 802.11b Wireless LAN (WLAN) operating at 10 Mbps and with fixed size frames. Assume we are operating without using RTS/CTS. Ignore the effect of propagation delay for the WLAN. Recall that each frame incurs the overhead of DIFS, Collision Window, ACK, and SIFS. Assume the combined overhead of DIFS, ACK, and SIFS to be 200 μs. Further assume that the average Collision Window size is 1.5 times the transmission time of the frame. Give an expression for the WLAN efficiency. Find the range of fixed frame size where Ethernet is more efficient compared to WLAN.

Solution

Ethernet efficiency $= X/(X + 5 \cdot 25.6 \cdot 10^{-6})$

WLAN efficiency $= X/(X + 200 \cdot 10^{-6} + 1.5 \cdot X)$

Hence, for Ethernet to be more efficient, we must have $200 \cdot 10^{-6} + 1.5 \cdot X > 5 \cdot 25.6 \cdot 10^{-6}$

$\Rightarrow$ Ethernet is more efficient for all frame sizes!

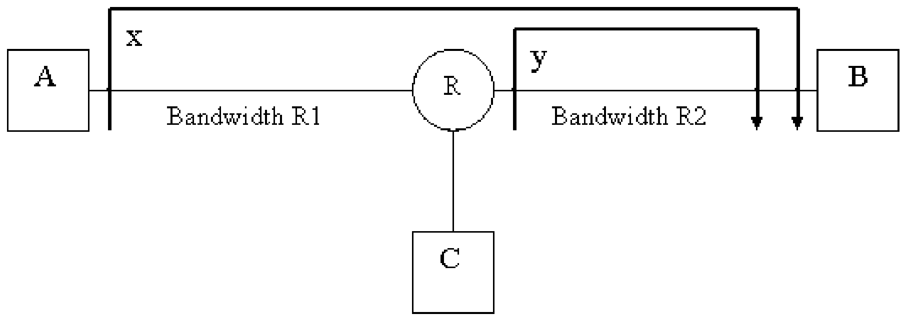

Figure 2: Diagram for Question 6.

Question 6 (13 points) Consider the configuration shown on Figure 2. Hosts A, B, and C are connected to each other via a router. The bandwidths of the links A-R and R-B are R_1 and R_2 respectively. The roundtrip time is the same between A-B, C-B and A-C.

(a) (8 points) Consider Figure 2. The rates for flows from A to B and C to B are x and y respectively. What are the values of x and y (as functions of R_1 and R_2) to which TCP converges? Figure 3 shows a diagram showing the sequence of rates (marked with solid lines) for specific values of R_1 and R_2. Explain the rate changes occuring according to AIMD in Figure 3. You may assume that the transmission times are negligible.

Solution

Generally, the two flows will share bandwidth equally on the R-B link. However, depending on the values of R_1 and R_2, this may or may not be possible. So, we look at it case by case:

1. $R_1 < R_2/2$: In this case, R_1 is the bottleneck for flow from A to B. Therefore, $x = R_1$ and $y = R_2 - x = R_2 - R_1$.
2. $R_1 \geq R_2/2$: In this case, bandwidth will be equally by both the flows on the R-B link. Therefore, $x = y = R_2/2$.

Now consider Figure 3. Some of the points have been labeled as A, B, C, D, E, F or G. The figure shows TCP behavior for a specific case of $R_1 < R_2/2$. As a result, the optimal point is $x = R_1, y = R_2 - R_1$, which has been marked with a solid circle. Point A is the starting point for the flows. Since $x + y < R_2$, additive increase keeps increasing x, y along the solid line AB (with slope 1) until it $x + y > R_2$ at B. So, at B, both the flows start dropping packets and multiplicative decrease (by a factor of 2) causes the rates to move to C. The rates then again undergo additive increase along the line CD until $x + y > R_2$ at D. Multiplicative decrease causes the rates to move to point E. Additive increase occurs again until we reach point F. At F, $x + y < R_2$ but $x > R_1$, therefore, x undergoes multiplicative decrease while y still keeps increasing.

So, we now reach point G. Additive increase from G causes us to reach D and then E, F, G, and so on. Thus, for this special case, the rates oscillate along two different lines!

(b) (5 points) Suppose another flow from A to C is also started (see Figure 4). Denote the rate of the additional flow as z. What are the values of x, y and z (as functions of R_1 and R_2) to which TCP converges?

Solution

The key point here is to realize that one of the links will be shared equally by the flows on that link. Now, we look at the two possible cases:

1. $R_1 < R_2$: In this case $x = z = R_1/2$. Therefore, $y = R_2 - x = R_2 - (R_1/2)$.
2. $R_1 \geq R_2$: In this case $x = y = R_2/2$. Therefore, $z = R_1 - x = R_1 - (R_2/2)$.

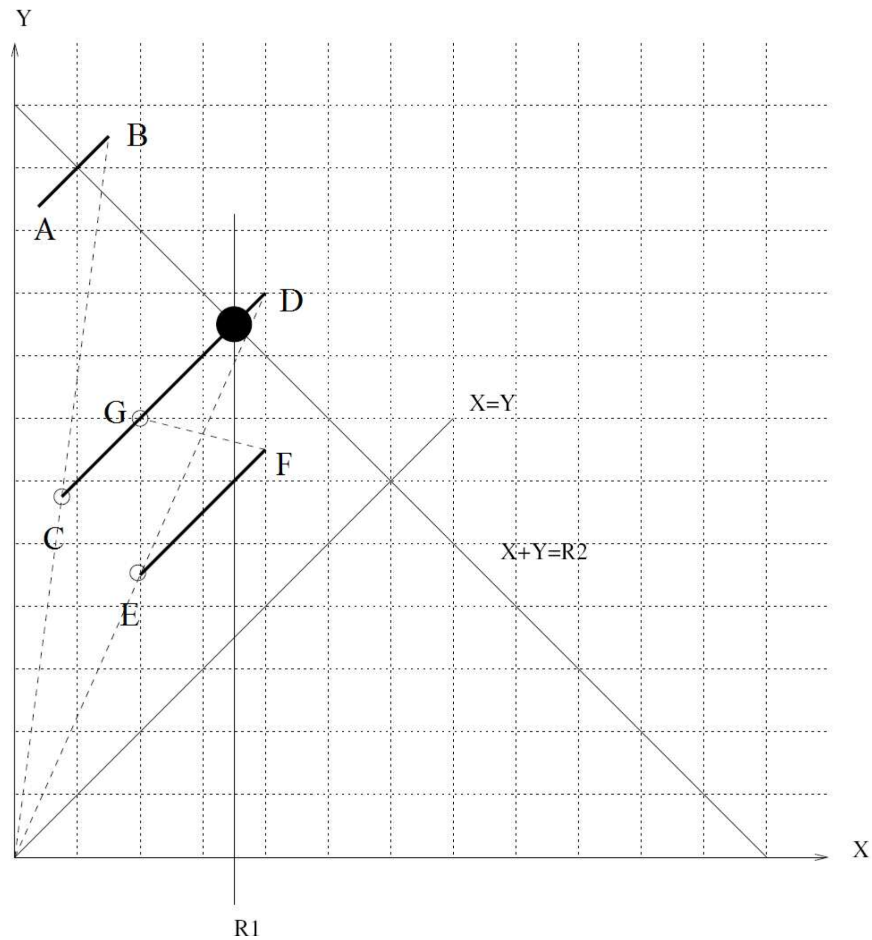

Figure 3: Diagram for Question 6(a). This shows the sequence of rates for specific values of R_1 and R_2. The points on the solid lines are the rates (x and y).

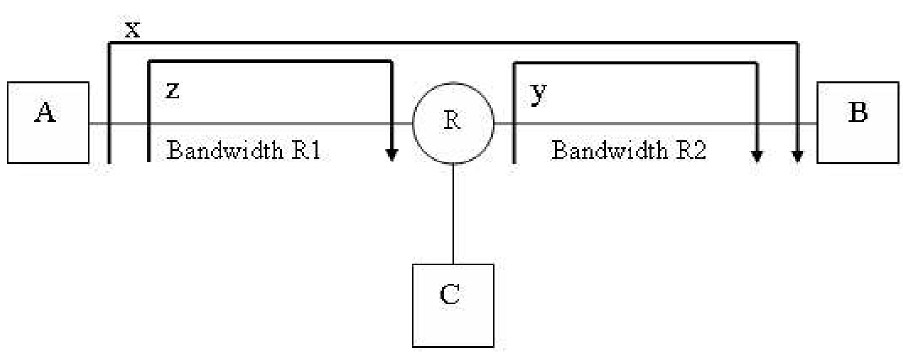

Figure 4: Diagram for Question 6(b).

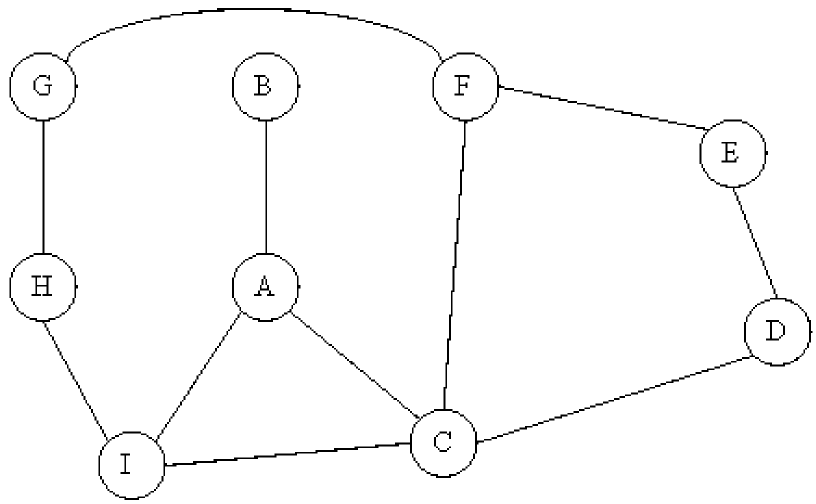

Figure 5: Diagram for Question 7.

Question 7 (10 points) Consider the network configuration shown in Figure 5. Assume that each link has the same cost.

(a) (5 points) Run Bellman-Ford algorithm on this network to compute the routing table for the node A. Show A's distances to all other nodes at each step.

Solution

The following table shows the routing table for A at each step until convergence.

A	B	C	D	E	F	G	H	I
0	1	1	∞	∞	∞	∞	∞	1
0	1	1	2	∞	2	∞	2	1
0	1	1	2	3	2	3	2	1

(b) (5 points) Suppose the link A-B goes down. As a result, A advertises a distance of infinity to B. Describe in detail a scenario where C takes a long time to learn that B is unreachable.

Solution

A advertizes a distance of infinity to B, but, C and I advertize a distance of 2 to B. Depending on the ordering of these messages, this might happen: I upon knowing that B can be reached from C with distance 2, concludes that it can reach B with a distance of 3 (via C) and advertises this to A. A concludes that it can reach B with a distance of 4 (via I), and advertises this to C. C concludes that it can reach B with a distance of 5 (via A). This continues forever if the distances are unbounded. This is called the count-to-infinity problem.

Question 8 (5 points)

Your local telephone company is adding DSL service to its existing network, and they need your help! From their frequency spectrum they have allotted the frequency band from 200-800 kHz, allowing simultaneous voice communication at 300-3300hz.The signal to noise ratio is 30dB (to simplify your calculations, you can consider $\log_2(1001) = 10$).

(a) (3 points) What is the capacity of this link?

Solution

6 Mbps

(b) (2 points) Should they advertise service with this link rate to their customers? Why or why not?

Solution

No. This is the theoretical maximum capacity of the link. The rate that they will obtain will be lower.

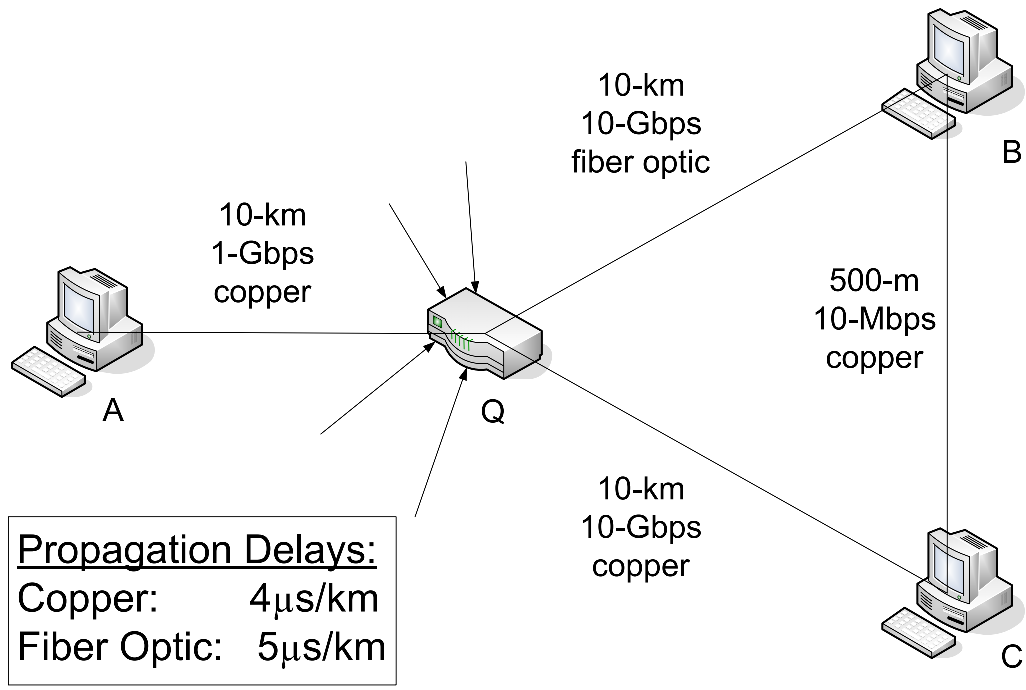

Figure 6: Diagram for Question 9.

Question 9 (10 points)

Consider the scenario shown in Figure 6. Host A is sending tiny packets to hosts B and C (neglect transmission time, just consider latency). Q is a store-and-forward switch with an average arrival rate of 10-Gbps and a buffer that contains 16 megabytes (base 10) of packets on average.

(a) (2 points) What is the average delay that packets will incur going through the switch?

Solution

Littles result: $16 \cdot 8 \cdot 10^6 = delay \cdot 10 \cdot 10^9$

$Delay = .0128s = 12.8ms = 12800\mu\text{s}$

(b) (2 points) What is the latency between A and B?

Solution

Prop of A-Q $= 10\text{km} \cdot 4\mu\text{s/km} = 40\mu\text{s}$

Q-B $= 10\text{km} \cdot 5\mu\text{s/km} = 50\mu\text{s}$

Latency $= 40\mu\text{s} + 50\mu\text{s} + 12800\mu\text{s} = 12890\mu\text{s}$

(c) (2 points) What is the latency between A and C?

Solution

Prop of A-Q $= 10\text{km} \cdot 4\mu\text{s/km} = 40\mu\text{s}$

Q-C $= 10\text{km} \cdot 4\mu\text{s/km} = 40\mu\text{s}$

Latency $= 40\mu\text{s} + 40\mu\text{s} + 12800\mu\text{s} = 12880\mu\text{s}$

(d) (2 points) What is the latency between B and C?

Solution

Prop of B-C $= 0.5\text{km} \cdot 4\mu\text{s/km} = 2\mu\text{s}$

Latency $= 2\mu\text{s}$

(e) (2 points) Without changing the network, can you propose a solution to decrease the delay between A and B?

Solution

The A-B delay can be decreased if you follow the route from A-C and then go from C-B.

$12880\mu\text{s} + 2\mu\text{s} = 12882\mu\text{s} < 12890\mu\text{s}$

Question 10 (9 points)

Below is the DNS record for a fictitious corporation, OK Computer:

Name	Type	Value	TTL (seconds)
okcomputer.com	A	164.32.15.98	86400 (1 day)
okcomputer.com	NS	thom.yorke.net	86400
okcomputer.com	NS	karma.okcomputer.com	86400
okcomputer.com	MX	android.okcomputer.com	60
lucky.okcomputer.com	A	164.32.12.8	86400
www.okcomputer.com	CNAME	lucky.okcomputer.com	86400
android.okcomputer.com	A	164.32.15.99	86400

(a) (3 points) If you type http://www.okcomputer.com into your web browser, to which IP address will your web browser connect?

Solution

www is a CNAME for lucky, whose IP is 164.32.12.8

(b) (3 points) If you send an e-mail to thom@okcomputer.com, to which IP address will the message get delivered?

Solution

The mail exchanger for okcomputer.com is android, whose IP is 164.32.15.99

(c) (3 points) The TTL field refers to the maximum amount of time a DNS server can cache the record. Give a rationale for why most of the TTLs were chosen to be 86400 seconds (1 day) instead of a shorter or a longer time, and why the MX record was chosen to have a 60-second TTL?

Solution

The TTL field tells the DNS server how long it may cache the record. Generally, DNS records do not frequently change, so a fairly large TTL is possible. However, if a record does change, the domain could be inaccessible for up to TTL time. 86400 seconds, or one day, is a good tradeoff. Okcomputer.coms administrators decided that they could handle one day of downtime if their IPs change. The MX record is required for delivery of e-mail. They do not want any messages lost if their IP changes, so one day of downtime is unacceptable. As a result, the record expires after only 60 seconds, and is thus resolved almost every time an email is sent.