CSE 473 – Introduction to Computer Networks	*Roch Guérin*
Quizz 3 Solution	
Your Name:	*10/02/2014*

1. (5 points total) The diagram below shows a typical DNS and http exchange, when *host* seeks to access a web page on *shop1.foo.com.*

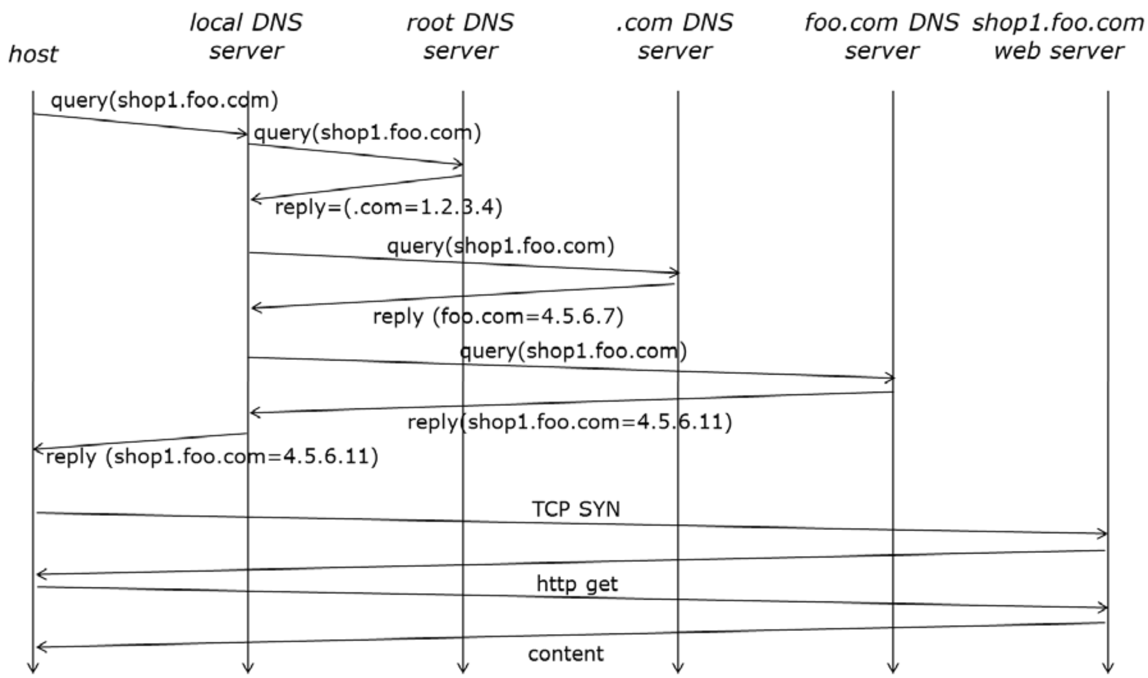

(2 points) Assuming all caches are empty at *host* and all DNS servers, what entries are present in the cache of the local DNS server at the end of the exchange?

.com => *1.2.3.4*
foo.com => *4.5.6.7*
shop1.foo.com => *4.5.6.11*

(1 point) What entries are present in the local DNS cache of *host*?

shop1.foo.com => *4.5.6.11*

(2 points) Assume that the user at *host* next clicks on a link that points to *pay.foo.com.* How many DNS server will be involved to handle this new query?

Two DNS servers will be involved. The local DNS server which receives the original query for pay.foo.com, and the DNS server for foo.com whose address was cached by the local DNS server that can therefore contact it directly to resolve the name pay.foo.com.

2. (5 points total) Consider a connection between hosts A and B that uses the Go-back-N protocol and uses packets of size 1.5 kbytes. The RTT (including propagation and queueing delays) is 2*20 ms = 40 ms (delay from A to B and back), and the path throughput is 100Mbits/sec (10^8 bits/sec).

   (2 points) What is the minimum window size (in bytes) that would allow continuous transmissions between A and B at a rate of 100 Mbits/sec in the absence of packet errors and losses?

   After A is done transmitting its first packet, it takes 40 ms for the ACK for that packet to come back. If A transmits continuously during that period, it would transmit a total of 0.04x10^8 bits = 4 Mbits or 500 kBytes. When added to the first packet that A transmitted, A would need a window size of at least 501.5 kBytes to be able to transmit continuously in the absence of packet losses and achieve the path maximum throughput of 100 Mbits/sec.

   (2 points) Assume now that A transmits maximum size packets, that the protocol's time-out is set to 2xRTT, *i.e.*, 80 ms (assume for simplicity that each packet has its own timer), and that the packet loss probability on the path between A and B is $p = 10^{-3}$. Under those assumptions, what is the average transmission throughput realized by the connection?

   The packet transmission time is equal to t_{pkt} = 120 μsec, so that the time to retransmit a lost packet is about 667.667 times the transmission time of a single packet. Based on the formula in the slides, we know that given a packet loss probability of $p = 10^{-3}$, the average time it takes to successfully transmit a packet is $T_{succ} = t_{pkt}$(1+666.667*0.001)/0.999 = 200.2 μsec. As a result, the connection transmits 1.5 kbytes every 200.2 μsec, which translates into a throughput of about 59.94 Mbits/sec.*

   (1 point) Under the same assumptions as in the previous question, consider now that the packet with sequence number 50 is transmitted at time t = 100ms and is lost. When will **packet 51** be delivered to the application at the receiver B, assuming no further packet losses?

   The loss of packet 50 is detected 80 ms after its transmission, i.e., at time 180 ms. This triggers its retransmission as well as that of all subsequently transmitted packets, including packet 51. Packet 51 leaves A at time 180 ms + 240 μsec = 180.24 ms, and arrives at B 20 ms later at time 200.24 ms, at which point it is delivered to the application