CS 640: Introduction to Computer Networks

Homework 3

Handed out: 11/27/2007, in class
Due: 12/6/2006, in class

Please type out or neatly print your answers. Add as much detail as you deem fit to your answers. For full credit you must explain how you obtained the answer and not just print the final result. Collaboration is discouraged strongly. Please try to work on your own. If you are stuck, you can discuss with your fellow students, but you are encouraged to talk to the course personnel first.

1. DNS using "dig"

Most Linux machines provide the **dig** program that allows you to query Domain Name Service (DNS) servers around the Internet. For more information on how to use dig, consult the man page.

When running dig for the purpose of this question, you should use the following format:

dig +norecurse @name.of.dns.server record-type domain-name

- **name.of.dns.server** is the hostname of the DNS server you wish to query.
- **record-type** is the type of DNS record you wish to retrieve, such as NS, ANY, MX, HINFO, A, or SOA.
- **domain-name** is the name of the host or domain you seek information on.

The DNS is a distributed architecture that uses hierarchical delegation. At the top of the system are the root name servers, who know which DNS server is responsible for each of the top-level domains (such as.com, .net, and .edu). If you send a root server a query for a particular machine, you will receive a reply listing the servers that have been delegated authority for those top-level domains, and you can recursively ask those servers to resolve the name.

a) To discover an actual chain of delegation, run a series of NS queries for WWW.CS.WISC.EDU. You may start with any of the root servers, and you should continue your sequence of queries until you stop getting new delegations (in some domains, this is indicated by a DNS server returning you a delegation pointing to itself, and in other domains this is indicated by a DNS server returning you a SOA record instead).

   Here is an example:

   Delegation chain for: AOL.COM:

Server queried	**NS delegations to**
A.ROOT-SERVERS.NET	A.GTLD-SERVERS.NET, G.GTLD SERVERS.NET
G.GTLD-SERVERS.NET	dns-01.ns.aol.com, dns-02.ns.aol.com, dns-06.ns.aol.com, dns-07.ns.aol.com
dns-02.ns.aol.com	dns-01.ns.aol.com, dns-02.ns.aol.com, dns-06.ns.aol.com, dns-07.ns.aol.com

This was produced by running the following commands:

dig +norecurse @a.root-servers.net NS aol.com
dig +norecurse @G.GTLD-SERVERS.NET NS aol.com
dig +norecurse @dns-02.ns.aol.com NS aol.com

Generate the delegation chain for WWW.CS.WISC.EDU. Present your results in a table like the one above. Each NS query will typically return two or more answers; choose among them at random. If you query a server and get a timeout, choose an alternate server.

b) The DNS is also used to translate IP addresses into hostnames. Again, the database is distributed in a hierarchical fashion, with a wrinkle. The most specific part of a domain name is on the left (i.e. WWW in WWW.CS.WISC.EDU), but the reverse is true of IP addresses (i.e., in 128.105.7.31, 128 is toplevel, 128.105 belongs to WISC.EDU, and 128.105.7 belongs to CS.WISC.EDU). Thus, address-to-name mapping is handled by reversing the bytes of the IP address and making queries in a special domain. To turn 128.2.198.101 into a hostname, various servers are sent queries seeking PTR records for 101.198.2.128.in-addr.arpa. The first query would be:

dig @a.root-servers.net PTR 101.192.3.128.in-addr.arpa

You will know you're done when your query gives you back a PTR record in addition to (or instead of) NS records. Fill in a table like the one above showing a query chain for the IP address 128.105.175.44. What is the DNS name corresponding to this IP address? What is the DNS name corresponding to the address 128.105.6.15?

2. Learning TCP using Ethereal

Answer the following questions about *tcp_pcattcp_n1.cap* network trace using Ethereal as the network traffic analyzer. Here the machine that initiates a TCP connection is referred to as the client.

(a). Which packets contain the three-way handshake in tcp_pcattcp_n1.cap? (Use packet numbers as shown in the trace.)

(b). What is the sequence number of the TCP SYN segment that is used to initiate the TCP connection between the client and the server? What is it in the segment that identifies the segment as a SYN segment?

(c). What is the sequence number of the SYNACK segment sent by server to the client computer in reply to the SYN? What is the value of the ACKnowledgement field in the SYNACK segment? How did the server determine that value? What is it in the segment that identifies the segment as a SYNACK segment?

(d). How much data does the receiver typically acknowledge in an ACK? Can you identify cases where the receiver is ACKing every other received segment. Which packets does packet 12 acknowledge? Answer the same for packet 14?

Answer the following questions about the trace httpWebBrowsing.cap using Ethereal as the network traffic analyzer.

(e). How many HTTP requests are sent to www.gnu.org? Which objects are requested? How big is each object?

(f). Is it persistent or non-persistent HTTP? How did you figure it out?

(g). Compute the average *response time* for www.google.com and www.gnu.org. Response time for a HTTP request is the time between the transmission of the first byte of the first request and the receipt of the last byte of the last HTTP object.

(h). Which packets contain the DNS request to translate www.gnu.org and the reply? What transport protocol are they using? What is the IP address of the server to which these packets are being sent? Which packet contains the DNS reply – i.e. the IP address of gnu.org? What is the IP address of www.gnu.org?

3. CDN

A browser on system A accesses index.html at www.cnn.com. Suppose CNN use Akamai to distribute its images and there are two images in the index.html named A.gif and B.gif. Akamai uses TTLs of 1 day, 30 minutes and 30 seconds for high-level DNS server NS records, low-level DNS server NS records and host Address records, respectively. The local DNS server for the system A is called D. Assume that the browser's cache and DNS cache on D are empty.

List the sequence of HTTP and DNS requests issued by A, D as well as Akamai's DNS and replica servers for:

(a) The first time A requests index.html from CNN,

(b) A new request issued by A for index.html 10s after the first request

(c) Another request issued 15 minutes after the first request and

(d) Another issued 45 minutes after the first request.

4. Multimedia

Video applications typically run over UDP rather than TCP because they cannot tolerate retransmission delays. However, this means video applications are not constrained by TCP's congestion control algorithms.

(a). What impact does this have on TCP traffic? Be specific about the consequences.

Fortunately, these video applications often use RTP, which results in RTCP "receiver reports" being sent from the sink back to the source. These reports are sent periodically (e.g. once a second) and include the percentage of packets successfully received in the last reporting period.

(b). Describe how the source might use this information to adjust its rate in a TCP compatible way.

5. TCP Congestion Control

(a). Calculate the buffering needed to to ensure that TCP fully utilizes a link if:

(1) TCP were to change its multiplicative decrease to decrease the window 3/4th the size after a loss.

(2) TCP were to change its additive increase to 2 packets per RTT

(b). Louis Reasoner has just finished his Internet Draft where he claims to have discovered three strategies that will eliminate congestion as a problem in the network. Each of these strategies would be very expensive to implement and create a logistical nightmare to coordinate. Even ignoring those issues, each proposal has a fatal flaw. Either it does not relieve congestion under all traffic patterns, or it creates another problem in the process of relieving congestion.

For each of these myths, explain why it does not relieve congestion or describe the new problem it creates. Draw figures of network topology, if needed, to illustrate your example.

(1). Since the throughput of TCP decreases when a packet is lost, and packets are lost because buffers in the network routers fill up, we will fix congestion by increasing the size of the buffers in the network by several orders of magnitude.

(2). Since buffers fill up because packets are waiting to go onto a link, we will increase the speed of *all* the links in the network by a factor of 10.

(3). To avoid congestion in the network, we will require access links to be slower than the links in the core of the network.

6. Fair Queuing and Weighted-Fair-Queuing (Attempt this question after Dec 4th lecture)

(a). Consider 10 flows with arrival rates of 1, 2, ..., 10 Mbps that traverse a link of 50Mbps. Compute the max-min fair share on this link. What is the fair share if the link capacity is 60 Mbps?

(b). Two flows A and B arrive at a router with a WFQ scheduling policy. The WFQ scheduling is modeled after GPS. Flow A has reserved 1/3 of the bandwidth on the outgoing link. Flow B has reversed 2/3 of the bandwidth on the outgoing link. Flow A's packets are one third the size of flow B's packets. What are the first 6 packets to leave the link?