COMP 431 - Problem Set 4

Joe Puccio

March 2, 2015

Collaborators: Fred Landis, Max Daum, Spencer Byers, Brandon Lam.

1.

a) Assuming that the latency provided was round-trip: we must solve $(1 - hit)(15 + 1400 + \frac{75 \cdot .05}{1.5} + \frac{75 \cdot .05}{10}) + hit(15 + \frac{75 \cdot .05}{10}) = 450$ for hit. We achieve a required hit ratio of $hit = 0.69$.

b) We can use the intensity for a link to estimate whether queueing delays should be expected. The intensity for a given link is simply the average number of bits being transmitted on the link over the transmission capacity of the link, so for the 1.5Mbps link, the intensity is $I = (.31)(75 \times .05)/1.5 = .775$ (at this point, we could potentially say that intensity is $> .6$ and thus queueing delays are not negligible, but we will continue). This intensity means on average the number of packets in the queue when a new packet arrives is $q = .775/(1 - .775) = 3.44$. A typical TCP packet is going to be ≈ 1500 bytes, which means 12000 bits. So when a request arrives, it will need to wait for the transmission of $3.44 \cdot 12000 = 41280$ bits, which would take the 1.5Mbps link 27.5 milliseconds. This is not really a negligible delay, as it is the same order of magnitude as the internal network's latency.

2.

a) The homepage will transfer in $50 + 50 + .2 = 100.2$ms. Then 4 TCP handshakes will commence simultaneously. Assuming the linked objects are downloaded in order and that one may have multiple simultaneous TCP connections with the same server (simply specify a different source port), meow-add.jpg and puss-chow-add.jpg will finish transferring at the same time, first, out of the top four files after 20.1ms. At this point, the two remaining TCP handshakes will begin, and because both of these files will finish after any of the original 4 and because the longest one takes 100.1ms to complete, the total time is $100.2 + 20.1 + 100.1 = 220.4$ms.

b) Assuming here that the subdomain www2 for dog.com is being hosted on a different server from dog.com (it is possible to host it on the same server) and that these requests are pipelined (all requests for the same server are sent in succession), we simply must add the one set of TCP connection times for each of the (4) unique servers, and then add the transmission time of every file which will also include roundtrip time (order does not matter because nothing is done in parallel). We achieve 310.8ms.

c) Parallel HTTP connections are still a good idea because the wait time that parallel HTTP connections really cut down on are ping delays (propagation delays) by overlapping these propagation delays to multiple servers. So long as your network interface card has a buffer (it does) to handle multiple streams of incoming data, significant time can still be saved by overlapping the propagation delays. This is because, in practice, these internetwork communications are usually going to take a lot longer than any sort of sorting through of interleaved packets in a buffer to get responses.

3.

a) Assuming that all 'www's are non-canonical and that www2.dogmyth.com is a different (load-balancing) server from dogmyth.com, the following records would be in the local DNS server's cache (structured so that each line is a new response to make part b easier):
$< .com, NS >< .com, A >$
$< www.dogmyth.com, CNAME >< dogmyth.com, A >$
$< www.ads.dogmyth.com, CNAME >< ads.dogmyth.com, A >$
$< .biz, NS >< .biz, A >$
$< cat.biz, NS >< cat.biz, A >$
$< www.ads.cat.biz, CNAME >< ads.cat.biz, A >$
$< bird.biz, NS >< bird.biz, A >$
$< www.ads.bird.biz, CNAME >< ads.bird.biz, A >$
$< www.cat.biz, CNAME >$
$< www2.dogmyth.com, A >$

b) Assuming that query/reply pairs are being counted, the local DNS server exchanges 10 query/reply protocol messages. This is because all CNAME records come with the IP (A record) of the canonical name, and because any time a NS record is returned, the A record for that nameserver is also returned so that the DNS server actually has an address at which to reach the nameserver. (Also note that many nameserver systems are not setup like this in practice, but rather one server acts as a nameserver for many hosts and their subdomains).

4.

We have that the total time that it takes for the file to be transferred is $t = \frac{S+40}{C} + \frac{F/S(S+40)}{C}$. We can find the minimum value by taking this function's derivative and setting it equal to 0, so we have the derivative $\frac{1-\frac{40F}{S^2}}{C} = 0$, and we find that this equation is zero when $F = \frac{S^2}{40}$, and so $S = \sqrt{40F}$. Now, we can verify that this is a minimum (rather than a maximum), by plugging in slightly larger and smaller values into the original equation, and we find that this indeed is a minimum.