HW #0 Solutions

ELEN E6761 - Computer Communication Networks
Fall 2000 Prof. Rubenstein

1. On my netscape browser, I enter in "google.com" and the google web-page appears. Describe the network-related events that occurred (where and when lookups took place, protocols used, etc.)

   Answer: The netscape application assumes this is an http request (note that http was not specifically specified). Then, the network address that corresponds to google.com must be determined. A DNS query is sent to the local name server, which responds with the IP address if it contains a record in its database. Otherwise (and assuming the query requests recursion and that all DNS servers support recursion), the local name server makes a similar request to its root name server, which either makes a request to the intermediate or authoritative name server if it does not have the IP address information. At some point, the information is obtained and passed back through a DNS reply - the information in the reply may be cached at the intermediate DNS servers upon return.

   Once the IP address has been identified, the http transfer can commence. Assuming HTTP/1.1 with pipelined, persistent connections, a single TCP connection is constructed containing and is sent to the identified IP address with an HTTP ascii message requesting http://google.com. The google server transmits the main page, while the client netscape browser identifies links within the page that require immediate display and within the same TCP connection (assuming pipelined, persistent connections), requests these additional links as well. Upon receiving all the data, the TCP connection is closed and the received information is displayed within the netscape window.

2. A company has opened a new office in three buildings that are adjacent to one another, and wishes to establish network connectivity with the Internet. How should these buildings be wired and connected?

   Answer: One possible solution is to install a LAN in each of the buildings, each LAN's hub connects to a single ethernet switch, which is connected to a router to the external Internet. Another solution is to bridge the three hubs in the three buildings together.

3. A file needs to be transferred reliably (i.e., no data loss) at a rate of 50 Kbs. The code has already been written such that the data is pumped into a socket at a rate of 50 Kbs. What problems can arise if the socket uses

   (a) TCP (socket protocol option SOCK_STREAM)?

   Answer: TCP will transfer the data reliably, but TCP regulates the transfer rate in order to avoid congesting the network. Hence, there is no guarantee that the data will be sent at a rate of 50 Kbs.

   (b) UDP (socket protocol option SOCK_DGRAM)?

   Answer: UDP does not regulate the rate at which data is transferred, but does not protect against packet loss or packet reordering.

4. Consider a set of IP routers that use the same routing protocol.

   (a) Give the name of a routing protocol.

   Answer: Some possibilities: DVRP (Distance Vector), OSPF (Open Shortest Path First)

(b) Describe in general how routes are computed (in enough detail that it distinguishes the protocol from other possible routing protocols).

Answer: For DVRP: each router maintains a vector of all other routers within the domain (i.e., the set of routers that use the same protocol). Each vector entry contains two entries: a measure of "distance" from itself to each of those routers, and the router that is the next hop on the path to the destination router. Periodically, each router informs its neighboring routers about its distances to all other routers in the network. Upon receiving this information, if a router can find a shorter distance to a given destination router, it updates its vector to include the new path.

(c) What state does an IP router maintain in order to forward an IP packet toward its destination (assume classless routing [CIDR])?

Answer: It must maintain enough information so that it can map each IP address to an outgoing interface. Typically, it stores the largest (longest) prefixes that cover sets of addresses that are forwarded on the same outgoing interface.

(d) A packet may traverse a path whose route contains routers running different routing protocols. How is this accomplished within the Internet?

Answer: Domains communicate via the Border Gateway Protocol (BGP). Routers at the "gateway" that interface to routers within another domain advertise the sequence of domains that a packet will take to reach a given prefix.

5. A sender, s, wishes to deliver a packet to $n+1$ receivers. It transmits the packet to a specially designated receiver, r_0, which then transmits the packet to each of the remaining n receivers, $r_1, \cdots, r_n$. The packet transmission from s to r_0 is successful with probability p_0. Transmission from r_0 to any $r_i, 1 \leq i \leq n$ is independently successful with probability p_1. Let $X_i, 0 \leq i \leq n$, be a random variable that equals 1 if r_i receives the packet, and equals 0 if r_i does not. Let $S = \sum_{i=0}^{n} X_i$. Calculate

(a) $\Pr(X_1 = 1)$
Answer: Both r_0, and then r_1 must receive the packet. Since the probability of receipt at the various points are independent, $\Pr(X_1 = 1) = p_0 p_1$

(b) $E[S]$
Answer: $E[S] = E[\sum_{i=0}^{n} X_i] = \sum_{i=0}^{n} E[X_i] = p_0 + np_0p_1$

(c) $\Pr(S = k)$ for $k < n + 1$

Answer: For $k = 0$, it must be the case that r_0 loses the packet, and if it does, no other receivers receive the packet. For $k > 0$, r_0 must receive the packet, and of the remaining receivers, $k - 1$ must receive the packet and $n - k + 1$ must lose the packet. It follows:

$$\Pr(S = k) \quad = \quad \begin{cases} 1 - p_0 & k = 0 \\ p_0 \binom{n}{k-1} {p_1}^{k-1} (1 - p_1)^{n-k+1} & k > 0 \end{cases}$$

(d) $\Pr(X_1 = 1 | X_0 = 1)$
Answer: Since r_0 receives the packet, it is simply the probability that the transmission from r_0 to r_1 is successful, and so $\Pr(X_1 = 1 | X_0 = 1) = p_1$

.

(e) $\Pr(X_1 = 1|X_2 = 0)$
Answer: Here, we use the fact that $\Pr(X_1 = 1|X_2 = 0) = \Pr(X_1 = 1, X_2 = 0)/\Pr(X_2 = 0)$. For the numerator, this is the probability that r_0 receives the packet and the (independent) transmission from r_0 to r_1 is successful, and that the (independent) transmission for r_0 to r_2 fails: $\Pr(X_1 = 1, X_2 = 0) = p_0 p_1(1-p_1)$. For the denominator, either the transmission to r_0 fails, or the transmission to r_0 succeeds and the transmission from r_0 to r_2 fails: $\Pr(X_2 = 0) = (1-p_0) + p_0(1-p_1)$. This yields:

$$\Pr(X_1 = 1|X_2 = 0) \quad = \quad \frac{p_0 p_1 (1-p_1)}{1 - p_0 p_1}$$

6. A sender, s, wishes to deliver m packets to a set of n receivers labeled $r_1, \cdots, r_n$ over a series of m rounds. Packet $j, 1 \leq j \leq m$ is sent to (but not necessarily received by) all receivers in the jth round. The transmission of any packet from s to any r_i is independently successful with probability p. Let R_j be a random variable that equals the number of rounds that must be performed so that all receivers receive *at least* j packets. Calculate

(a) $\Pr(R_j = j)$ for $1 \leq j \leq m$
Answer: The first j transmissions must be received by all n receivers. Since each transmission is independent of all others (to other receivers, as well as for different packets to the same receiver), we have $\Pr(R_j = j) = p^{nj}$

(b) $\Pr(R_1 > i)$ for $1 \leq i \leq m$

Answer: This is the probability some receiver has failed to receive a packet by the ith round. This is $1 - \Pr(R_1 \leq i)$, i.e., 1 minus the probability that all receivers have received at least one packet by the ith round. Since the event of a successful packet transmission is independent of the success or failure of all transmissions, define $X_{i,j}$ to be a r.v. that equals 1 if and only if r_j receives at least one packet in i rounds, and 0 otherwise. Then $\Pr(R_1 \leq i) = \prod_{j=1}^{n} \Pr(X_{i,j} = 0)$. For any j, the probability that r_j gets at least one packet is the 1 minus the probability that it gets no packets, hence for any j, $\Pr(X_{i,j} = 0) = 1 - (1-p)^i$. Substituting through, we get

$$\Pr(R_1 > i) \quad = \quad 1 - (1 - (1-p)^i)^n$$

(c) $E[R_1]$ (express as a sum)
Answer: For a random variable, X, that takes only integer values, $E[X]^{\infty} i \Pr(X = i) = i = 0 i = 0^{\infty} \Pr(X > i)$. For this problem, since $\Pr(R_1 > m) = 0$, we have

$$\begin{aligned} E[R_1] &= \sum_{i=0}^{m-1} \Pr(R_1 > i) \\ &= \sum_{i=0}^{m-1} \left(1 - (1 - (1-p)^i)^n\right) \end{aligned}$$

(d) $\Pr(R_j > j + 2)$ for $1 \leq j \leq m - 2$.
Answer: Define $Y_{i,k}$ to be a random variable that equals the number of packets that r_k receives after i rounds. Then

$$
\begin{aligned}
\Pr(R_j > j+2) &= 1 - \prod_{k=1}^{n} \Pr(Y_{j+2,k} \leq j+2) \\
\Pr(Y_{j+2,k} \leq j+2) &= \Pr(Y_{j+2,k} = j+2) + \Pr(Y_{j+2,k} = j+1) + \Pr(Y_{j+2,k} = j) \\
&= p^{j+2} + (j+2)(1-p)p^{j+1} + \binom{j+2}{2}(1-p)^2 p^j
\end{aligned}
$$

Hence, $\Pr(R_j > j+2) = 1 - \left(p^{j+2} + (j+2)(1-p)p^{j+1} + \binom{j+2}{2}(1-p)^2 p^j\right)^n$

7. Station A transmits data on a communication channel in bursts. Let T_i be a random variable that equals the start-time of A's ith transmission burst, and let L_i be a random variable that equals the length of time of A's ith transmission. The L_i are exponentially distributed with rate λ_1 (i.e., $\Pr(L_i > t) = e^{-\lambda_1 t}$). After A finishes its ith transmission, it waits a period of time that is exponentially distributed with rate λ_2.

   (a) If $\lambda_1 = \lambda_2$, then what is $\Pr(T_{i+1} - T_i > t)$ (in terms of λ_1 and t)?
   Answer: We make use of the fact that exponentially distributed random variables are memoryless. The probability to be computed is simply the probability that two events occur in time less thatn t, where each event is exponentially distributed with rate λ_1.
   We have that $\Pr(L_i > t) = e^{-\lambda_1 t}$, by taking the derivative with respect to t, we get that $dens(L_i > t) = -\lambda_1 e^{-\lambda_1 t}$. Define D_i to be the random variable that equals the off-time after burst L_i. Then $\Pr(D_i > t) = e^{-\lambda_1 t}$ (i.e., it has the same distribution as L_i, and

$$
\begin{aligned}
\Pr(T_{i+1} - T_i > t) &= \int_{x=0}^{t} dens(L_i = x) \Pr(D_i > t - x) dx + \Pr(L_i > t) \\
&= \int_{x=0}^{t} -\lambda_1 e^{-\lambda_1 x} e^{-\lambda_1 (t-x)} dx + e^{-\lambda_1 t}
\end{aligned}
$$

   Working through the integral, we get

$$
\Pr(T_{i+1} - T_i > t) = (1 + \lambda_1 t) e^{-\lambda_1 t}
$$

   (b) Let x be the fraction of time that the channel is used by A. What is λ_2 in terms of λ_1 and x?

   Answer: Note that the r.v.s representing the length of times of the various bursts and outages are independent. Hence, we can treat the system as a renewal process, such that $x = E[L_i]/(E[T_{i+1} - E[T_i] = \lambda_1/(\lambda_1 + \lambda_2)$. Via algebraic manipulation, we get $\lambda_2 = \lambda_1(1-x)/x$

   (c) Let t_1 be the time that A starts transmitting its first burst, and assume we know that at time $t_2 > t_1$, it is still transmitting that burst. What is $\Pr(L_1 > t | L_1 > t_2 - t_1)$? (NOTE THE CHANGE IN THE QUESTION: originally asked what is $\Pr(L_i > t | L_i > t_2 - t_1)$)

   Answer: Again, we take advantage of the memoryless property of the exponentially distributed r.v. We know that at time t_2 that the transmission is still active, hence we are asking what the probability is that it extends for another time of $t - t_2$, hence $\Pr(L_i > t | L_i > t_2 - t_1) = e^{-\lambda_1 (t - t_2)}$