Midterm Exam for Computer Networks (CNT 6215)

Fall 2012

>>> SOLUTIONS <<<

Welcome to the Midterm Exam for *Computer Networks*. Read each problem carefully. There are eight required problems (each worth 12 points – you get 4 points for correctly following these instructions). There is also an additional extra credit question worth 10 points. You may have with you a calculator, pencils and/or pens, erasers, blank paper, and one 8.5 x 11 inch "formula sheet". On this formula sheet you may have anything you want (definitions, formulas, homework answers, old exam answers, etc.) as **handwritten by you in pencil or ink** on both sides of the sheet. Photocopies, scans, or computer generated and/or printed text are not allowed on this sheet. Note to tablet (iPad, etc.) users – you may **not** print-out your handwritten text for the formula sheet. You have 75 minutes for this exam. **Please use a separate sheet of paper for the answer to each question**. Good luck and be sure to show your work!

Problem #1

Each sub-problem worth 3 points.

Answer the following questions regarding the basics principles and concepts of networks.

a) Sketch the 5-layer Internet protocol model – identify the key protocols for each layer. In one sentence, or less, describe the function (or purpose) of each layer.

```
  +--------------+
5 | Application  | -- User application
  +--------------+
4 |     TCP      | -- Provides guaranteed delivery of data
  +--------------+
3 |      IP      | -- Responsible for moving packets from end to end
  +--------------+
2 |     Link     | -- Provides point-to-point data transport
  +--------------+
1 |   Physical   | -- Transmission of bit stream
  +--------------+
```

b) Define *protocol* and *interface*. Give one example of each.

```
Protocol = Complete set of rules regarding information exchange between same level
layers between sites. Interface = Complete set of rules regarding information exchange
between adjacent layers within a site. Examples: Protocol = HTTP, Interface = Sockets
```

c) In 1988 Clark defined the top level goal of the Internet as "effective technique for multiplexed utilization of existing interconnected networks." What were the second level goals for the Internet as defined by Clark?

```
The second level goals were 1) communications continue despite loss of networks or
gateways, 2) support multiple types of communication service, 3) accommodate a variety
of networks, 4) permit distributed management, 5) be cost effective, 6) easy host
attachment, and 7) resource use accountable.
```

d) What is the end-to-end principle as defined by Saltzer? Given an example.

```
The end-to-end principle states that if a required function can only be implemented in
an end node, then implementing it in the network is not possible (but, may be useful
for performance reasons). For example, error detection and correction can never be
fully done by the network; it must be done in the end node in any case.
```

Problem #2

Each sockets function is 1 pt. The connect and send/recv flows are 1 pt each.

Sockets is the standard API between Internet applications and the TCP (or UDP) protocol. Sockets is based on a client/server model. Sketch the flowcharts (giving the key sockets function calls) for a streams-based server and client. The flowcharts should show data exchange between the server and client.

```
        socket()
           |
           V
         bind()
           |
           V
        listen()                  socket()
           |                         |
           |                         V
           |<------------------ connect()
           V                         |
        accept()                     |
           |                         |
           V                         V
    send() / recv() <-----> send() / recv()
           |                         |
           V                         V
      closesocket()            closesocket()
```

Problem #3

SW formula is 4 pts. Set-up is 6 pts. Calculation is 2 pts.

Consider a host with a 100 Mb/s interface. The host sends 1250 byte packets on a link using a sliding window protocol. If the maximum allowable window size is 10, what is the maximum possible length link for which this host can achieve full 100 Mb/s throughput (that is, 100% utilization of the link)? Assume that signal propagation is 5 microseconds per mile. Assume that the link is error-free.

The key formula for sliding window flow control is:

$$U_{SW} = \max\left(1, \frac{W \cdot t_{fr}}{2 \cdot t_{pr} + t_{fr}}\right)$$

We can write

$$1.0 = \frac{10 \cdot \left(\frac{1250 \cdot 8}{100 \cdot 10^6}\right)}{2 \cdot t_{pr} + \left(\frac{1250 \cdot 8}{100 \cdot 10^6}\right)}$$ and solve for t_{pr} = 4.5 x 10^{-4}.

For propagation of 5 microsecond per mile we get $\frac{4.5 \cdot 10^{-4}}{5 \cdot 10^{-6}} = 90$ miles

For case (2) 100% utilization of the link can never be achieved with a sliding window (or any other ARQ) protocol. Think about it – any loss (which must occur due to the definition of the link being not error-free) will trigger a retransmission of at least the lost packet and thus utilization is now less than 100%.

Problem #4 Each sub-problem worth 6 points.

Answer the following short questions regarding Routing.

a) List the criteria for a good routing algorithm (and for its implementation as a protocol).

A good routing algorithm/protocol should be fast, simple, not generate much network traffic (so, low overhead), not create loops, be stable, and converge to an optimum solution (lowest cost route).

b) How can routing instability occur in a network? Give an example. What can be done to "fix it"?

Routing instability can occur when a dynamic metric (such as lowest delay) is used for the cost metric. Oscillation between two parallel paths (call them A and B) can occur when one path (say, A) has lower delay all packets will try to take this path, then it's delay will increase over the delay of path B and now all packets will try to take path B, path A now having lower delay, and the process repeats (oscillates). Dampening (that is, making the responsiveness to cost metric relative "slow") can be used to minimize oscillation. This is a difficult problem that requires control theoretic insights.

Problem #5 Use judgement (see paper).

In an analysis of Ethernet (CSMA/CD protocol), Metcalfe came-up with a result for link efficiency as

$$E = \frac{(P/C)}{(P/C) + W \cdot T}$$

Where P is the number of bits in a packet, C is the peak channel capacity (in b/s), T is the slot time (defined as the worst case collision time), and W is the "mean number of slots of waiting in a contention interval before a successful acquisition of the Ether by a station's transmission". Derive this result (**hint:** you need to derive W).

Assume that there a Q stations continuously queued to transmit a packet and that each of these stations has probability $1/Q$ of transmitting in a given slot. Let A be the probability that exactly one stations transmits successfully in a given slot,

$$A = Q \cdot (1/Q) \cdot (1 - (1/Q))^{Q-1} = (1 - (1/Q))^{Q-1} .$$

That is, there are Q ways in which one station can transmit in this slot while Q – 1 stations do not attempt to transmit in this slot. A is geometric, from which the mean waiting time is them,

$$W = \frac{1 - A}{A} .$$

Problem #6 Each sub-problem worth 3 points.

Answer the following questions regarding performance modeling and experimental design.

a) What is a model (precisely define it), what is the goal of building a model, and why do we build models (and not just study actual systems)?

"A model is a representation (physical, logical, or functional) that mimics another object under study" (Molloy 1989). The goal of a building a model is to be able to understand the system under study – to be able to predict its performance. We build models because they are cheaper, easier, faster, and/or safer to experiment with compared to actual systems.

b) What is experimentation? What is experimental design? What is the goal of design of experiments?

Experimentation is a method of investigating causal relationships among variables. Experimental design is the specification of the combinations of factors, factor levels, and number of replications. The goal of design of experiments is to determine the maximum amount of information about a system with the least amount of effort.

c) Consider a system with two factors A and B where each factor has two levels (A1, A2, B1, and B2) and a single response variable (which is a function of the factor levels). Show an example where the factors are non-interacting. Show an example where the factors are interacting. Two factors interact if the effect of one depends on the level of the other.

	A1	A2
B1	3	5
B2	6	8

Not-interacting – As A increases by 2 the response variable increases by 2 independent of the level of B

	A1	A2
B1	3	5
B2	6	9

Interacting – As A increases by 2 the response variable increases by 2 or 3 dependent on the level of B

d) What is the advantage (and disadvantage) of a full-factorial experiment design compared to a simple experiment design?

Advantage: A full-factorial experiment design can find all possible factor interactions where a simple experiment design may miss such interactions. Disadvantage: A full-factorial design has more experiments to run that a simple design.

Problem #7

Each sub-problem worth 4 points.

Answer the following questions regarding Markov models.

a) Precisely define the Markov property? What is a DTMC? A CTMC?

Past history is completely summarized by the specification of a current state. A DTMC is a Discrete Time Markov chain for system with discrete (periodic) time transitions – for each transition a probability of transition is given. A CTMC is a Continuous Time Markov Chain with continuous time transitions (i.e., transitions can occur at any time) – for each transition a rate of transition is given.

b) What is a P matrix? What is a Q matrix? Briefly describe the properties of each.

A P matrix is the matrix showing all state transition probabilities for a DMTC. A Q matrix is the matrix showing all state transition rates for a CMTC. For a P matrix each row sums to 1. For a Q matrix each row sums to 0 with each diagonal element being the negative sum of all other (non-diagonal) elements in the row.

c) Convert the below Q matrix to a P matrix and then solve it (for steady state probabilities).

$$Q = \begin{bmatrix} -6 & 6 \\ 2 & -2 \end{bmatrix}$$

We can use uniformization, $P = I + (1/\gamma) \cdot Q$ where γ is the maximum absolute value in Q (6 in the case of the above example). So, we get

$$P = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \left(\frac{1}{6}\right) \cdot \begin{bmatrix} -6 & 6 \\ 2 & -2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 1/3 & 2/3 \end{bmatrix}$$

From $\pi = \pi \cdot P$ and $\sum \pi = 1$ we can write $\pi_0 = (1/3) \cdot \pi_1$ and $1 = \pi_0 + \pi_1$, which solves to $\pi_0 = 1/4$ and $\pi_1 = 3/4$.

Problem #8

4 pts for CTMC, 4 pts for Q matrix, 2 pts for set-up, 2 pts for correct solution.

Assume that two people in an office make calls at a rate λ calls per minute and stay on the call for a mean 1/μ minutes. Assume that the time between calls and the time to stay on a call are exponentially distributed, and that calls are independent. Let λ = 1 and μ = 2. What is the probability of 0, 1, or 2 phone lines being busy?

As our first step, we draw the CTMC.

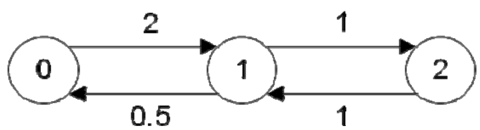

We then write the Q matrix:

$$Q = \begin{bmatrix} -2 & 2 & 0 \\ 0.5 & -1.5 & 1 \\ 0 & 1 & -1 \end{bmatrix}$$

We then write the three equations in three unknowns and solve for the steady state probabilities as:

$$-2\pi_0 + 0.5\pi_1 = 0$$
$$2\pi_0 - 1.5\pi_1 + 1\pi_2 = 0$$
$$\pi_0 + \pi_1 + \pi_2 = 1$$

Note that one of the equations is $\pi_0 + \pi_1 + \pi_2 = 1$ to break the linear dependence of the equations. We solve and get $\pi_0 = 1/9, \pi_1 = 4/9$, and $\pi_2 = 4/9$.

Extra Credit

+5 for the assumption, +5 for the significance.

What is the “Kleinrock Independence Assumption” as described in Kleinrock’s “On the Modeling and Analysis of Computer Networks” *(Proceedings of the IEEE,* August 1993)? What is the significance of this assumption?

“The Independence Assumption assumes that the length of a message is chosen independently from the exponential distribution each time it enters a switching node in the computer network!” The significance of this assumption is that it makes analysis of a network of queues (that is, modeling switching nodes in a computer network) trivial by allowing the application of the M/M/1 delay formula independently to each queue.

Humor

From: http://www.physics.ucf.edu/~costas/COURSES/ExamCartoon.jpg

I hope that everyone did well ☺