

>>> **SOLUTIONS** <<<


Welcome to exam #1 in *Capacity Planning* (CIS 4930/6930). You have 75 minutes. Read each problem carefully. There are six required problems (each worth 16 points - you get 4 points for "free") and one extra credit problem worth 5 points. You may have with you a calculator, pencils, erasers, blank paper, lucky rabbit's foot, and one 8.5 x 11 inch "formula sheet". On this formula sheet you may have anything you want (definitions, formulas, etc.) *handwritten by you*. You may use both sides. Computer generated text, photocopies, and scans are not allowed on this sheet. Please submit your formula sheet with your exam. Please start each numbered problem on a new sheet of paper and do not write on the back of the sheets (I really do not care about saving paper!). Submit everything in problem order. No sharing of calculators. Good luck and be sure to show your work!

Problem #1 (10 minutes) - (a) = 6 pts, (b), (c) = 5 pts each

a) Describe the capacity planning process (i.e., what are the inputs and output to and from the process). Hint: A figure may be helpful to your description.

Capacity planning is a human process that takes workload, system parameters, and desired service (of a system) as input and determines the saturation point and system alternatives.

b) Describe with a graph typical system behavior with respect to offered load and response time.



Non-linear behavior with a knee and an asymptote.

c) When are experiments on a model more appropriate than experiments on a real system? Give at least three reasons.

Use a model when the real system is too expensive or too dangerous to experiment with. Also, use a model when the real system does not yet exist.

Problem #2 (10 minutes) - (a), (b), (c), (d) = 4 pts each

a) What is the primary determinant of performance in a computer system?

The large difference between memory and CPU speed, and how this difference is handled, is the primary determinant of performance in a computer system

b) What is round trip time (RTT) in a network?

RTT is the time it takes for a packet to be sent from a sender to a receiver and then returned to the sender from the receiver.

c) Assume a type of network packet that is 53 bytes in length of which 5 bytes are header. What is the maximum utilization possible for a network that uses these types of packets?

$$U = 48 / (48 + 5) = 90.6\%$$

d) On a 10-Mbps Ethernet link, what is the transmission time for a 1500 byte packet? What is the approximate propagation delay if the link is 1000 feet in length?

$T_{\text{xmit}} = (1500 * 8) / 10e6 = 1.2$ milliseconds. 1000 feet is about 1 microsecond in speed-of-light propagation.

Problem #3 (15 minutes) - (a), (b) = 3 pts each, (c), (d), (e) = 2 pts each, (f) = 4 pts

a) Define probability. Why is probability theory of interest to capacity planners?

Probability can be experimentally defined as $\text{Pr}[\text{event}] = \lim_{n \rightarrow \infty} (\# \text{ observed events in } n \text{ trials} / n)$. Probability is of interest because the behavior of most computer systems, networks, and so on is random. Probability theory is the toolkit to understand systems with random behavior.

b) What is a random variable (define it)?

A random variable is a function that maps a real number to every possible outcome in the sample space. A random variable can be continuous or discrete.

c) What is a probability density function (pdf/pmf)? What is a probability distribution function (PDF/CDF)? Carefully describe each and give the key properties of each.

The density function is simply the plot of the value of a RV. The area under a density function sums to 1.0. The distribution function is the running sum of the density function. Thus a distribution function is 0 at $-\infty$, 1.0 at ∞ , and is monotonically increasing.

d) For a fair 6-sided die, what is the probability of rolling a 6 followed by a 5 followed by a 4 (i.e., three rolls of the die with a 6 appearing first, then a 5 on the next roll, and then a 4 on the next roll)? Show your work. Name the “rule” that applies to your calculation.

Independent events, so... $1/6 \times 1/6 \times 1/6 = 1 / 216$.

e) For a fair 6-side die, what is the probability of rolling a 6 or a 5 or a 4 on one roll of the die? Show your work. Name the “rule” that applies to your calculation.

Mutually exclusive events, so... $1/6 + 1/6 + 1/6 = 1/2$.

f) Give the formulas for arithmetic mean and harmonic mean. Describe when (i.e., for what kinds of measurements) each should be used.

$\bar{X}_A = \frac{1}{N} \sum_{i=1}^N X_i$ and $\bar{X}_H = \frac{N}{\sum_{i=1}^N \frac{1}{X_i}}$. Use arithmetic means for times and harmonic mean for rates.

Problem #4 (15 minutes) - (a), (b) = 4 pts each, (c) = 8 pts

Here are some measurements (e.g., of web server response time in milliseconds)...


110, 120, 140, 120, 120, 100, 110, 100, 140, 120

a) Compute the mean, variance, standard deviation, and coefficient of variation of this population of measurements

$E[X] = \frac{1}{10} (110 + 120 + 140 + 120 + 120 + 100 + 110 + 100 + 140 + 120) = 118$

$$\sigma^2 = \frac{1}{10}((110-118)^2 + (120-118)^2 + \dots + (120-118)^2) = 176$$

b) Plot a histogram for the measurements.



c) Find the 95% confidence interval for this set of measurements. A T-score table is given at the end of this exam.

$$s = \sqrt{\frac{1}{10-1}((110-118)^2 + (120-118)^2 + \dots + (120-118)^2)} = 13.98$$

For 9 degrees of freedom, $t = 2.26$ (from the second column in the table for 95% confidence interval), so $H = 2.26 \cdot \left(\frac{13.98}{\sqrt{10}}\right) = 9.99$

and thus the population mean lies between 108 and 128 with 95% confidence.

Problem #5 (15 minutes) - (a), (b) = 8 pts each -- syntax is graded only very lightly

a) Write a C-code function to compute and return the mean of the values stored in an array X of length N. The function prototype is:

```
double mean(double X, unsigned int N);
```

```
double mean(double X, unsigned int N)
{
    double mom1;
    int i;

    for (i=0; i<N; i++)
        mom1 = mom1 + (1.0 / N) * X[i];
    return(mom1);
}
```

where X is the array of values of length N. The return from the function is the mean.

b) Write a C-code function to measure and return the elapsed time for how long it takes to add-up all the value in an array X of length N. The function prototype is:

```
double timed_add(double X, unsigned int N);
```

```
double timed_add(double X, unsigned int N)
{
    double sum, elapsed;
    struct timeb start, stop;
    int i;

    ftime(&start);
    for (i=0; i<N; i++)
        sum = sum + X[i];
    ftime(&stop);
    elapsed = ((double) stop.time + ((double) stop.millitm * 0.001)) -
              ((double) start.time + ((double) start.millitm * 0.001));
    return(elapsed);
}
```

where X is the array of values of length N. The return from the function is the elapsed time in seconds.


Problem #6 (10 minutes) - (a), (b) = 5 pts each, (c) = 6 pts

a) Name a tool that can be used to measure and graph CPU utilization for Windows NT/2K.

Perfmon or taskmanager.

b) What is SNMP? Identify and sketch the key components of SNMP. Name a tool that can be used to GET and view SNMP variables.

SNMP is Simple Network Management Protocol. The key components of SNMP are a manager and a remote device containing an agent and MIB (Management Information Base). The agent communicates with the manager to, for example, send variables from the MIB to the manager. The MIB contains variables which are typically statistical counters (e.g., bytes out for an interface). SNMPTOOL (MRTG also) can be used to GET and view SNMP variables.



c) What is MRTG? How, roughly speaking, does MRTG work?

MRTG is Multi-Router Traffic Grapher. MRTG uses a Perl script to GET SNMP variables (typically, bytes in and out of a router port) and then plots them on a graph. The graph is inserted into a webpage which is stored in a preset directory. If MRTG is run on a webserver, the web pages are then viewable from anywhere. A key feature of MRTG is its ability to compress graphs from a minute, to hourly to daily to weekly and so on views. This compression (achieved by blocking) makes archiving of data not require ever increasing storage space.

Extra Credit problem: (??? minutes) - 5 pts

Assume that you have made measurements over several months on network utilization for a large corporate network. These measurements were made every 10 minutes. The autocorrelation of these measurements shows a spike (i.e., a value close to 1) for a lag of 144 measurement samples, but is roughly zero for all other lag values great than about 6. What can you speculate is occurring?

144 x 10 is 1440 minutes or 24 hours. Such a daily spike could be caused by a cron job for daily back-up. A back-up would consume a large spike in utilization for few to 10's of minutes.

T-scores. Selected values of $t_{\alpha/2; N-1}$

	$\alpha/2 = 0.05$	$\alpha/2 = 0.025$
N - 1	t	t
4	2.13	2.78
5	2.02	2.57
6	1.94	2.45
7	1.90	2.37
8	1.86	2.31
9	1.83	2.26
10	1.81	2.23
11	1.80	2.20
12	1.78	2.18