Systems Q Exam

January, 1998

This exam has five questions and takes 90 minutes to complete. Use a bluebook for both scratch paper and exam answers, and clearly label the answers to the questions. If you feel a question is ambiguous, raise your hand and a proctor will help you. If you still feel the question is ambiguous after talking to a proctor, state reasonable assumptions and solve the problem.

1. [Basic Knowledge]

For each of the following scenarios, provide a quantitative answer within a factor of 5 of the exact answer (explain briefly how you derive the value if you are unsure about the 5x margin). Assume the system in question is a typical 1997 desk-top computing system:

(a) capacity of "hard drive"

(b) length of time to get a block from the "hard drive"

(c) size of main memory

(d) length of time for an email message to go across campus

(e) length of time for an email message to go across the US on a "good day"

(f) Take two hosts. How long does it take one to "ping" the other (i.e. send a small message round-trip)

(i) if they are connected by an Ethernet.

(ii) if they are connected by an 28.8K baud modem using PPP or SLIP.

(iii) if they are across the US connected by the Internet on a "good day." Assume the hosts have an Ethernet connection to the Internet.

(g) Describe the storage hierarchy of a modern PC or workstation. For each level, provide the typical storage size and access time.

(h) What is the memory bandwidth necessary to drive a high-resolution (1024x768, 24 bits/pixel) color display?

2. [Concurrent Programming]

(a) Many processors provide so-called "interlock instructions" to facilitate solving the mutual exclusion problem. An example of such an instruction is the test-and-set instruction

```
TS(v,p):  <  p := v;    v :=  true  >
```

where the notation <S> signifies that S is performed indivisibly.

Using TS, complete the following program fragment for ensuring that execution of the code fragment CS is mutually exclusive with respect to other similarly protected code fragments (i.e., add code at the point labeled XXX). Assume processes are named 1, 2, ... N.

```
Variables: s, p_1, p_2, ... p_N : Boolean
Initially: s=false, p_i = false for all i.
```

Variable s is shared among processes, while each of the variables p_i are local to process i. In process i:

```
entry_i:  p_i := true
          XXX
CS
exit_i:   s := false
```

(b) The fetch-and-add instruction FA(w,p,e) is an interlock instruction that has been proposed in connection with shared-memory massively-parallel multiprocessors. It's effect is given by:

```
FA(w, p, e):  <  p := w;  w := w+e >
```

Consider the following purported solution to the mutual exclusion protocol:

```
Variables: s, p_1, p_2, ... p_N : Integer
Initially: s=1, p_i = 0 for all i.
```

In process i:

```
entry_i:  FA(s, p_i, -1)
          while p_i <= 0 do
              a_i: FA(s, p_i, 1)
              b_i: FA(s, p_i, -1)
              end
CS
exit_i:   FA(s, p_i, 1)
```

There are problems with this protocol. What are they? Give specific execution scenarios that illustrate the problems (a_i and b_i are labels that may be useful in your answer).

(c) Speculate about the advantages of fetch-and-add over TS on multiprocessors.

3. [Operating Systems]

(a) Describe **both** the sequence of events **and** the software subsystems involved in a Remote Procedure Call between two user level processes running on Unix workstations attached to a shared Ethernet hub. Draw pictures to illustrate your answer.

(b) Estimate the times required for each of the subsystems (in milliseconds)

(c) If you were asked to improve the performance of cross-machine RPC, what hardware and software changes would you make? For each change, explain the exact change you would make, and describe why you would make that change. Remember that you have the freedom to completely redo every component of the operating system and the networking subsystem.

4. [Architecture]

Most first generation RISC processors had a simple 5-stage integer pipeline. Typically, the function of the stages were:

1. IFETCH - access the instruction cache at the PC and fetch instruction
2. OFETCH - decode instruction and fetch source registers
3. EXECUTE - perform arithmetic operation on operands
4. MEMORY - access data cache for load/store instructions
5. WBACK - write result into destination register

Under ideal steady-state operation, a new instruction enters the pipeline each clock cycle and takes one cycle in each stage. For each of the following instruction sequences,

(i) describe the problem caused by the above pipelining,
(ii) describe (or name) the most reasonable solutions for first generation RISC processors, and
(iii) explain the impact/cost of each solution.

Keep your answers brief. Just "hitting the keywords" (if you know them) is enough.

(a)

```
add     r2 <- r3, r4      // r2 := r3 & r4
andi    r2 <- r2, 0x7f    // r2 := r2 & 0x7f
ori     r2 <- r2, 0x80    // r2 := r2 | 0x80
```

(b)

```
cmp     r6, r7
beq     else
addi    r1 <- r1, 1
else:
```

(c)

```
ld      r6 <- 0(r5)       // r6 := memory[r5]
addi    r6 <- r6, 1       // r6 := r6 + 1
addi    r5 <- r5, 4       // r5 := r5 + 4
```

5. [Databases]

(a) The structure of a nested-loops (NL) join algorithm between relations R and S look is:

```
for all data r in R {
    for all data s in S {
        try to join r and s
    }
}
```

Your database system uses the operating system virtual memory (with LRU replacement) to buffer disk resident data. You wish to measure the performance of a NL join between R (fixed size) and S (variable size). For different sizes of relation S along the X-axis, you plot the total number of page I/Os on the Y-axis of a graph (access to a page already in memory does not count as an I/O).

i) why would there be a "step" in the plotted function, and what does it correspond to?

ii) if you could change the replacement policy, is it possible to smooth the step?

(b) You want to see Michael Jordan play for the Bulls against the Knick's on Jan 31st. Although tickets are tough to find, you manage to reach Ticket Master (telephone ticketing agents) and buy a ticket. The agent charges your credit card, enters stuff in their database system, issues you a ticket number, and tells you that the transaction is "committed". Right after this, however, Ticket Master's entire computer system crashes. Do you expect to find someone else in your seat on the day of the game? (explain your answer in terms of the actions of the database system, not the actions of crazy basketball fans).

(c) Suppose you were to build a web search engine using the following tools:

- a web crawler that retrieves HTML documents
- an HTML parser that extracts the non-formatting words in an HTML document,
- an object-relational database that has support for strings, and URLs as data types.

i) Suggest a schema for database relations that will hold the data needed to satisfy search requests. (hint: minimally, you need to record which words occur in which document)

ii) Write an SQL query that could be used to service a typical (Alta Vista style) search request, or comes close to it.

iii) Give an interesting query over this data that can be posed in SQL but not in a keyword search (e.g., a Web search engine such as Alta Vista).