Computer Science Comprehensive Exam

Architecture

February 5, 1996

This is a closed-book exam. Answer all of the questions. Be concise in your answers. You will receive partial credit for partially correct answers, but extraneous remarks may count against you.

1. Under a two's complement representation for integers, a left shift can replace an integer multiplication by a power of 2. *Is a right shift the same as an integer division by a power of 2?*

2. Consider a memory hierarchy that has a 64 entry TLB and a 64Kbyte cache. The TLB is fully associative and uses a random replacement policy. The cache is direct mapped, physically addressed, and physically tagged. The cache block size is 16 bytes.

   A memory reference can suffer from three different types of misses: a TLB miss, a cache miss, and a page fault. *Is it possible for a memory reference to incur a TLB miss but not a cache miss?* If not, explain why not. If so, briefly sketch a program that incurs more TLB misses than cache misses.

3. Compare and contrast **hardwired** versus **microprogrammed** control. In particular, explain why you would choose one over the other.

4. The network interfaces in dedicated parallel machines, such as the IBM SP2, typically support both **programmed I/O** and **DMA I/O**. During the execution of a parallel program, a node, i.e., a processor, is dedicated to the execution of that program. No other program is allowed to execute on that node. Thus, the program can be safely allowed to access the network interface directly from user-level. Most message-passing libraries choose to use programmed I/O instead of DMA I/O because the time from when the message is sent by one program to when it is received by its partner, especially for small messages, is lower. *Why would programmed I/O have lower latency?*

5. On a SCSI bus, each I/O device negotiates with the host computer to determine the burst transfer rate to use—usually the fastest mutually supported speed. Assume that the maximum burst transfer rate supported by the host computer is 20 MByte/second. Furthermore, assume that all devices have adequate buffering so that they can maintain their sustained transfer rate even when competing for bus time.

   Assume a tape drive with a sustained transfer rate of 500 KBytes/second and a burst rate of 4 MBytes/second is attached to the SCSI. You want to attach disk drives to the same bus, each with a sustained transfer rate of 6 MBytes/second and a burst rate of 20 MBytes/second. *How many disk drives can you attach to the bus and expect to be able to run all devices simultaneously at full speed? What will the bus utilization be in that case?* (You may use 1K=1,000 and 1M=1,000,000 to simplify your calculations.)