

Enterprise-Scale Software Development (COMP5348)

Semester 1, 2009

Tutorial Week 5

Part A: Paper-based exercises

Q1. You have a 1GB file on an external disk. If the bandwidth obtainable when reading sequentially from the disk is 25MB/s, how long will it take to read the file (suppose that there is initial set-up time of 5ms, for disk seek etc)? Suppose instead that the file is read in a sequence of separate operations in random order, where each operation reads a 1KB block and requires its own set-up time. How long will it take to complete reading every block of the file?

Q2 Suppose requests arrive every 100ms, and that 90% of requests are of type A, and 10% of requests are of type B. Each type A request requires processing that executes 2M instructions (say that on average these instructions use 2 cycles each), and it also requires 10ms of disk activity; type B requests need no I/O but they need 100M instructions (many of these are expensive multiplications, and the average is 4 cycles per instruction). If the CPU speed is 2 GHz, what load is being placed on the CPU and the disk by these requests? If the workload arrival rate becomes faster, estimate when and why saturation might occur.

Q3. Suppose we can buy three different system configurations

Configuration X: 2 GHz CPU, 2GB of slow main memory, 128kB of cache; cost is \$1250

Configuration Y: 2 GHz CPU, 2GB of slow memory, 3MB of cache; cost is \$1500

Configuration Z: 2 GHz CPU, 2GB of slow memory, 64MB of cache; cost is \$6000

Suppose that excluding times for load/store, each instruction takes 2 cycles. The access time for cache is 10ns, and for slow main memory is 70ns.

Suppose that our workload does on average 1 load or store per 10 instructions, and it has locality properties so that the hit rate is 50% for the small cache (configuration X), 95% for configuration Y, and 100% for configuration Z. What is the average performance (in instructions/sec) for each configuration? What is the performance per dollar for each configuration?

Part B: Run perfmon on the lab machine, and explore the behaviour when opening a large file from disk (drive C) compared to opening it from across the network (from drive U). See <http://adminfoo.net/2007/04/windows-perfmon-top-ten-counters.html>