

CS451 Project #1

Benchmarking

Instructions:

- **Due date: 11:59PM on Monday, 02/03/14**
- **Maximum Points: 100%**
- **Maximum Extra Credit Points: 25%**
- *You should work in teams of 2 for this assignment.*
- *Please post your questions to the Piazza forum.*
- *Only a softcopy submission is required; it must be submitted to “Digital Drop Box” on Blackboard.*
- *For all programming assignments, please submit just the softcopy; please zip all files (report, source code, compilation scripts, and documentation) and submit it to BB.*
- *Name your file as this rule: “PROJ#_LASTNAME1_LASTNAME2.{zip|tar|pdf}”. E.g. “Proj1_Doe_Smith.tar”.*
- *Late submission will be penalized at 10% per day (beyond the 7-day late pass).*

1 Your Assignment

This project aims to teach you how to benchmark different parts of a computer system, from the CPU, GPU, memory, disk, and network.

You can be creative with this project. You are free to use any programming languages (C, C++, Java, etc) and any abstractions such as sockets, threads, events, etc. that might be needed. You are free to use any machines for your development as long as it will work under Linux for your final evaluation.

In this project, you need to design a benchmarking program that covers four of the five components listed below; **to get full credit, you need to implement just 4 components**; *the maximum extra credit you can get for this assignment is 25, although there are 50 extra possible points you could attempt:*

1. **CPU:** Measure the processor speed, in terms of floating point operations per second (Giga FLOPS, 10^9 FLOPS) and integer operations per second (Giga IOPS, 10^9 IOPS); measure the processor speed at varying levels of concurrency (1 thread, 2 threads & 4 threads) → $2*3 = 6$ experiments
 - a. *hint, modern processors can do multiple instructions per cycle, so make sure to give your benchmark good code to allow it to run multiple instructions concurrently*
 - b. *for 5% extra credit, vary the number of threads and find the optimal number of concurrency to get the best performance*

2. **GPU:** Measure the GPU speed, in terms of floating point operations per second (Giga FLOPS, 10^9 FLOPS) and integer operations per second (Giga IOPS, 10^9 IOPS); measure the processor speed with full concurrency (mapping 1 thread per core) → 2 experiments
 - a. *Hint: You will have to use either CUDA or OpenCL to do this assignment; timing might also be tricky, compared to measuring timing on the host system*
 - b. *Make sure to make your GPU code generic, and not hard coded to match your specific GPU; the TAs will run your code on a GPU that might be different than your GPU, and your code must still be correct*
 - c. *For 5% extra credit, measure the read and write memory bandwidth of the GPU memory with different size messages (1B, 1KB, 1MB)*
3. **Memory:** Measure the memory speed of your host; your parameter space should include read+write operations (e.g. memcpy), sequential access, random access, varying block sizes (1B, 1KB, 1MB), and varying the concurrency (1 thread & 2 threads). The metrics you should be measuring are throughput (Megabytes per second, MB/sec) and latency (milliseconds, ms) → $1*2*3*2 = 12$ experiments
 - a. *hint: you are unlikely going to be able to do this benchmark in Java, while C/C++ is a natural language to implement this benchmark*
 - b. *for 5% extra credit, vary the number of threads and find the optimal number of concurrency to get the best performance*
4. **Disk:** Measure the disk speed; your parameter space should include read operations, write operations, sequential access, random access, varying block sizes (1B, 1KB, 1MB, 1GB), and varying the concurrency (1 thread & 2 threads). The metrics you should be measuring are throughput (MB/sec) and latency (ms)
 - a. *Hint: there are multiple ways to read and write to disk, explore the different APIs, and pick the fastest one out of all them → $2*2*4*2 = 32$ experiments*
 - b. *for 5% extra credit, vary the number of threads and find the optimal number of concurrency to get the best performance*
5. **Network:** Measure the network speed in terms of bytes/second; your parameter space should include the loopback interface card (between 2 processes on the same node), the TCP protocol stack, UDP, varying packet/buffer size (1B, 1KB, 64KB), and varying the concurrency (1 thread & 2 threads). The metrics you should be measuring are throughput (Megabits per second, Mb/sec) and latency (ms) → $1*2*3*2 = 12$ experiments
 - a. *for 5% extra credit, run the same experiments over a 1Gb/s Ethernet switch between two different systems*

Other requirements:

- You must write all benchmarks from scratch. You can use well known benchmarking software to verify your results, but you must implement your own benchmarks. Do not use code you find online, as you will get 0 credit for this assignment.
- All of the benchmarks will have to evaluate concurrency performance; concurrency can be achieved using threads. Be aware of the thread synchronizing issues to avoid inconsistency or deadlock in your system.
- All of these benchmarks could be done on a single machine, but for some network tests, you could use 2 machines.
- Experiments should be done in such a way that they take multiple seconds to minutes to run, in order to amortize any startup costs of the experiments; that means that for some of the operations that are really small (e.g. 1B), you might need to do many thousands or even millions of them to run long enough to amortize the costs of the benchmark overheads.

- Not all timing functions have the same accuracy; you must find one that has at least 1ms accuracy or better.
- Since there are many experiments to run, find ways (e.g. scripts) to automate the performance evaluation.
- Make sure your machine is idle when running the benchmarks as it would help to improve reliability and consistency of the results.
- No GUIs are required. Simple command line interfaces are fine.

2 What you will submit

When you have finished implementing the complete assignment as described above, you should submit your solution to 'digital drop box' on blackboard. Each program must work correctly and be detailed in-line documented. You should hand in:

1. **Design Doc (10%):** A separate (typed) design document (named design.pdf or design.txt) of approximately 1-3 pages describing the overall program design, and design tradeoffs considered and made. Also describe possible improvements and extensions to your program (and sketch how they might be made).
2. **Manual (10%):** A detailed manual describing how the program works. The manual should be able to instruct users other than the developer to run the program step by step. The manual should contain example commands to invoke each of the four benchmarks.
3. **Source code (30%):** All of the source code; in order to get full credit for the source code, your code must have in-line documents, must compile, and must be able to run the sample benchmarks from #2 above. Please specify which student contributed to what code.
4. **Performance (50%):** Since this is an assignment aimed at teaching you about benchmarking, this is one of the most important part; you must evaluate the four benchmarks with the entire parameters space mentioned in Section 1, and put them in a separate file (named evaluation.pdf). You must produce graphs to showcase the results. Please combine data and plot on the same graph wherever possible, for either compactness reasons, or comparison reasons. Don't forget to plot the average and standard deviation, as opposed to just a simple value. Also, you need to explain each graph's results in words. Hint: graphs with no axis labels, legends, well defined units, and lines that all look the same, are likely very hard to read and understand graphs. You will be penalized if your graphs are not clear to understand. Please specify which student contributed to what benchmark experiments.

Please put all of the above into one .zip or .tar file, and upload it to 'digital drop box' on blackboard'. The name of .zip or .tar should follow this format: *PROG#_SURNAME_FIRSTNAME.{zip|tar}*. Please do NOT email your files to the professor and TA!! You do not have to submit any hard copy for this assignment, just a soft copy via Black Board.

Grades for late programs will be lowered 10% per day late beyond the 7-day late pass.