

263-2300-00: How To Write Fast Numerical Code

Assignment 1: 100 points

Due Date: Th, March 9th, 17:00

<http://www.inf.ethz.ch/personal/markusp/teaching/263-2300-ETH-spring17/course.html>

Questions: fastcode@lists.inf.ethz.ch

Submission instructions (read carefully):

- (Submission)
Homework is submitted through the Moodle system <https://moodle-app2.let.ethz.ch/course/view.php?id=3122>. Before submission, you must enroll in the Moodle course.
- (Late policy)
You have 3 late days, but can use at most 2 on one homework, meaning submit latest 48 hours after the due time. For example, submitting 1 hour late costs 1 late day. Note that each homework will be available for submission on the Moodle system 2 days after the deadline. However, if the accumulated time of the previous homework submissions exceeds 3 days, the homework will not count.
- (Formats)
If you use programs (such as MS-Word or Latex) to create your assignment, convert it to PDF and name it homework.pdf. When submitting more than one file, make sure you create a zip archive that contains all related files, and does not exceed 10 MB. Handwritten parts can be scanned and included or brought (in time) to Alen's, Georg's or Gagandeep's office. Late homeworks have to be submitted electronically by email to the fastcode mailing list.
- (Plots)
For plots/benchmarks, **provide (concise) necessary information for the experimental setup (e.g., compiler and flags) and always briefly discuss the plot and draw conclusions**. Follow (at least to a reasonable extent) the small guide to making plots from the lecture.
- (Code)
When compiling the final code, ensure that you use optimization flags. **Disable SSE/AVX for this exercise when compiling**. Under Visual Studio you will find it under Config / Code Generator / Enable Enhanced Instructions (should be off). With gcc their are several flags: use `-mno-abm` (check the flag), `-fno-tree-vectorize` should also do the job.
- (Neatness)
5% of the points in a homework are given for neatness.

Exercises:

1. (20 pts) Get to know your machine
Determine and create a table for the following microarchitectural parameters of your computer.
 - (a) Processor manufacturer, name, and number.
 - (b) Number of CPU logical and physical cores.
 - (c) CPU-core frequency.
 - (d) CPU maximum frequency. Does your CPU support Intel Turbo Boost Technology?
 - (e) Tick or tock model?

For one core and **without** considering SSE/AVX:

- (d) Latency/Throughput/Gap for floating point additions.
- (e) Latency/Throughput/Gap for floating point multiplications.
- (f) Latency/Throughput/Gap for the `rcp` instruction (if supported).
- (g) Latency/Throughput/Gap for fused multiply - add (FMA) operations (if supported).
- (h) Maximum theoretical floating point peak performance in both flop/cycle and Gflop/s.

Notes:

- Intel calls throughput what is in reality the gap = 1/throughput.
- The manufacturer's website will contain information about the on-chip details. (e.g. [Intel 64 and IA-32 Architectures Optimization Reference Manual](#)).
- On Unix/Linux systems, typing `cat /proc/cpuinfo` in a shell will give you enough information about your CPU for you to be able to find an appropriate manual for it on the manufacturer's website (typically AMD or Intel).
- For Windows 7/10 "Control Panel/System and Security/System" will show you your CPU, for more info "CPU-Z" will give a very detailed report on the machine configuration.
- For Mac OS X there is "MacCPUID".
- Throughout this course, we will consider the FMA instruction as two floating point operations.

2. (10 pts) Cost analysis

Consider the following algorithm for calculating the sum of first $2n$ terms in the $\cot^{-1}(\frac{1}{x})$ series (we assume $n > 1$):

```
1 double calc (int n, double x) {
2     double num1 = x, num2 = num1*x*x;
3     int den1 = 1, den2 = 3;
4     double sum = num1/(double)den1 - num2/(double)den2;
5     while (den1 < 4*n - 4){
6         num1 = num1 * x * x * x * x * x;
7         num2 = num1 * x * x;
8         den1 = den1 + 4;
9         den2 = den2 + 4;
10        sum = sum + num1/(double)den1 - num2/(double)den2;
11    }
12    return sum;
13 }
```

- (a) Define a suitable detailed floating point cost measure $C(n)$.
- (b) Compute the cost $C(n)$ of the function `calc`.

3. (25 pts) Matrix multiplication

In this exercise, we provide a C source [file](#) for multiplying two matrices of size n and a C header [file](#) that times matrix multiplication using different methods under Windows and Linux (for x86 compatibles).

- (a) Inspect and understand the code.
- (b) Determine the exact number of (floating point) additions and multiplications performed by the `compute()` function in `mmm.c`.
- (c) Using your computer, compile and run the code (Remember to turn off vectorization as explained on page 1!). Ensure you get consistent timings between timers and for at least two consecutive executions.
- (d) Then, for all square matrices of sizes n between 100 and 1500, in increments of 100, create a plot for the following quantities (one plot per quantity, so 3 plots total). n is on the x-axis and on the y-axis is, respectively,
 - i. Runtime (in cycles).
 - ii. Performance (in flops/cycle).
 - iii. Using the data from exercise 1, percentage of the peak performance (without vector instructions) reached.
- (e) Briefly discuss your plots.

4. (20 pts) Performance Analysis

Consider the function *babs*:

```
1 inline double babs (double x) {
2     union { uint64_t i; double d; } u = { .d = x };
3     u.i = u.i & 0x7FFFFFFFFFFFFFFF;
4     return u.d;
5 }
```

Assume that the elements of vectors x, y, u and z of length n are combined as follows:

$$z_i = z_i + x_i \cdot y_i + u_i \cdot y_i + x_i \cdot z_i + \text{babs}(u_i \cdot x_i).$$

- Write a C/C++ compute() function that performs the computation described above on arrays of doubles. Save the file as combine.c(pp).
- Within the same file create a benchmarking infrastructure based on the timing function that produces the most consistent results in Exercise 3.
- Then, for all two-power sizes $n = 2^4, \dots, 2^{22}$ create performance plot with n on the x-axis (choose logarithmic scale) and performance (in flops/cycle) on the y-axis. Create two series such that the first has all optimization flags disabled, and the second series has all optimizations flags enabled (except for vectorization). Randomly initialize all arrays. For all n repeat your measurements 30 times reporting the median in your plot.
- If you have an Intel processor, run the same tests again, such that you make sure that Intel Turbo Boost is disabled (or enable it if the previous plot was generated with Turbo Boost disabled).
- Briefly explain eventual performance variations in your plot and the effects of Turbo Boost.

5. (20 pts) Bounds

Consider the three artificial computations below. The functions operate on a input vector and store the results in an output array:

```
1 void artcomp1(float x[], float y[], int N) {
2     for (int i = 0; i < N; i++)
3         y[i] = x[i] * 1.3 + y[i];
4 }
5 void artcomp2(float x[], float y[], int N) {
6     float len = N/2; //assume 2 divides N
7     for (int i = 0; i < 2 * len; i += 2) {
8         y[i] = x[i] * 1.4;
9         y[i + 1] = x[i + 1] + 1.4;    }
10 }
11 void artcomp3(float x[], float y[], float z[], int N) {
12     float len = N/3; //assume 3 divides N
13     for (int i = 0; i < 3 * len; i += 3) {
14         y[i] = x[i] * z[i];
15         y[i + 1] = x[i + 1] * z[i+1];
16         y[i + 2] = x[i + 2] + 3.3;    }
17 }
```

We consider a Core i7 CPU based on a Haswell processor. As seen in the lecture, it offers FMA instructions (as part of AVX2) that compute $y = a * x + b$ on floating point numbers. Consider the information from the lecture slides on the throughput of the according operations. Assume that divisions are performed with the regular div on Port 0. Assume the bandwidths that are given in the additional material from the lecture: [Abstracted Microarchitecture](#)

- Determine the exact cost (in flops) of each function.
- Determine an asymptotic upper bound on the operational intensity (assuming empty caches and considering both reads and writes) of each function

- (c) Consider only one core and determine, for each function, a hard lower bound (not asymptotic) on the runtime (measured in cycles), based on:
- The op count. Assume that the code is compiled using `gcc` with the following flags: `-fno-tree-vectorize -mfma -march=core-avx2 -O3` and that FMAs are used as much as possible. Be aware that the lower bound is also affected by the available ports offered for the according computation. (see lecture slides)
 - Loads, for each of the following cases: All floating point data is L1-resident, L2-resident, L3-resident, and RAM-resident. Consider best case scenario (peak bandwidth).

How to disable Intel Turbo Boost

Intel Turbo Boost is a technology implemented by Intel in certain versions of its processors that enables the processor to run above its base operating frequency via dynamic control of the processor's clock rate. It is activated when the operating system requests the highest performance state of the processor.

BIOS

Intel Turbo Boost Technology is typically enabled by default. You can only disable and enable the technology through a switch in the BIOS. No other user controllable settings are available. Once enabled, Intel Turbo Boost Technology works automatically under operating system control. When access to BIOS is not available, few workarounds are possible:

Linux

Linux does not provide interface to disable Turbo Boost. One alternative, that works, is disabling Turbo Boost by writing into MSR registers. Assuming 2 cores, the following should work:

```
wrmsr -p0 0x1a0 0x4000850089
wrmsr -p1 0x1a0 0x4000850089
```

To enable it:

```
wrmsr -p0 0x1a0 0x850089
wrmsr -p1 0x1a0 0x850089
```

This method has been criticized [here](#) and, [here](#) stating that the OS can circumvent the MSR value, using opportunistic strategy. But so far in our tests, we have observed that Linux conforms to the MSR value. An alternative method would be to use `cpupower`, as explained in the [ArchLinux Wiki](#), as well as the `intel_pstate` driver. Unfortunately, we can not confirm deterministic behavior across different kernel versions with this method.

Mac OS X

Disabling Turbo Boost in OS X can be done easily with the [Turbo Boost Switcher for OS X](#). Note that the change is not persistent after restart. The method also writes to the MSR register, and shares the same weaknesses as the Linux approach.

Windows

Windows does not provide any functionality to disable Intel Turbo Boost. The only effective way of disabling is using the BIOS. On some Intel machines however, it is possible to fix the CPU multiplier such that the resulting frequency corresponds to the nominal frequency of the CPU. [ThrottleStop](#) provides this functionality with a convenient GUI. "Disable Turbo" will effectively fix the frequency such that it corresponds to a behaviour of a CPU with disabled Turbo Boost.