

# Data Processing on Modern Hardware

## Assignment 1


Handout: 29<sup>th</sup> April 2020

Due date: 6<sup>th</sup> May 2020


### Set-up

The main goal of this assignment is to measure the latency of each level of the memory hierarchy of your machine. To achieve that, you need to write a short micro-benchmark that accesses the system's memory with specific access patterns. To do that you need to create an array in main memory and initialize it appropriately. Then the elements of the array should be accessed in one of the following patterns:


1. sequential



2. sequential, but in reverse order



3. random (from the perspective of the hardware prefetcher).



Your micro-benchmark should take as input the following parameters:

1. the total size of the array  $S$ ,
2. the stride size  $d$ ,
3. the access pattern  $p$ .

## Tasks

- a Write a C program that implements such a micro-benchmark.
- b Run the benchmark for the three access patterns with different values for  $d$  and  $S$  on your system (e.g.,  $d \in [8B, \dots, 16kB]$  and  $S \in [4kB, \dots, 512MB]$ ) and plot the measured values.
- c Based on the measurements answer the following questions.
  - What observations do you make?
  - Does the experiment match your expectations / hypothesis?
  - How can you explain the observed behavior?
  - What conclusions can be drawn from your hardware based on the results?

Note that if you have access to different computers (with different architectures (e.g., AMD vs. Intel vs. ARM), or older vs. newer hardware, or different cache and memory configurations. etc.), you are welcome to experiment with them and report the obtained numbers.

## Submission guidelines

Please use your gitlab folder for this course to submit your solutions to the assignment. Create a sub-folder titled Assignment 1, which will include:

- Your micro-benchmark source code with a Makefile.
- The script used to run the experiments.
- A short report that answers the questions.

## Hint

In the article “*What Every Programmer Should Know About Memor*” by Ulrich Drepper, you can find a lot of information for implementing the micro-benchmark and interpreting the results.