



University of
Zurich^{UZH}

Department of Physics



Scientific Programming with Python

Hardware Speedup Exercises

June 25, 2020

General comments

- Do not forget to activate the virtual environment:
`source ~/school_venv/bin/activate.`
- Depending on your hardware you can increase the number of CPUs on your virtual machine.
- It might be interesting to run the examples on the host-computer in case you have a running python environment there.

Exercise 1: Optimizing arithmetic expressions

1. Use script `poly.py` to check how much time it takes to evaluate the next polynomial:

$$y = .25*x**3 + .75*x**2 - 1.5*x - 2$$

with x in the range $[-1, 1]$, and with 10 millions points.

- you can execute the script with different arguments. For example:
`./poly.py --library numpy --expression-index 0.`
- Set the `--library` argument to `numexpr` and take note of the speed-up versus the “numpy” case. Why do you think the speed-up is so large?

If you get a “Permission denied” error you need to set execution permission to the file with the following command: `chmod +x poly.py`

If your `python` command points to Python2, you need to change the first line (also for the other scripts) to `#!/usr/bin/env python3`.

2. The expression

$$y = ((.25*x + .75)*x - 1.5)*x - 2$$

represents the same polynomial than the original one, but with some interesting side-effects in efficiency. Repeat this computation (`--expression-index 1`) for `numpy` and `numexpr` and get your own conclusions.

- Why do you think numpy is doing much more efficiently with this new expression?
 - Why the speed-up in numexpr is not so high in comparison?
 - Why numexpr continues to be faster than numpy?
3. The C program `poly.c` does the same computation than above, but in pure C. Compile it like this:
- ```
gcc -O3 -o poly poly.c -lm
```
- and execute it with `./poly`
- Why do you think it is more efficient than the above approaches?

## Exercise 2: Evaluating transcendental functions

4. Evaluate the expression `sin(x)**2+cos(x)**2` in `poly.py`, a function that includes transcendental functions (`--expression-index 3`).
- Why the difference in time between NumPy and Numexpr is so small?
5. In `poly.c`, comment out expression 1) (around line 56) and uncomment expression 3) – the transcendental function). Don't forget to compile again.
- Do this pure C approaches go faster than the Python-based ones?
  - What would be needed to accelerate the computations?

## Exercise 3: Using Numba

The goal of Numba is to compile arbitrarily complex Python code on-the-fly and executing it for you. It is fast, although one should take in account the compile times.

6. Open `poly-numba.py` and look at how numba works.
- Run several expressions and determine which method is faster. What is the compilation time for numba and how it compares with the execution time?
  - Raise the amount of data points to 100 millions. What happens?

## Exercise 4: Parallelism

7. Be sure that you are on a multi-processor machine use the
- ```
y = ((.25*x + .75)*x - 1.5)*x - 2
```
- expression in `poly-mp.py` by using the argument `--expression-index 1`. Repeat the computation for both numpy and numexpr for a different number of processes (numpy) or threads (numexpr). Pass the desired number with `--threads` to the script.

- How does the efficiency scale?
 - Why do you think it scales that way?
 - How is the performance compared with the pure C computation?
8. With the previous examples, compute the expression:
 $y = x$
That is, do a simple copy of the 'x' vector. What is the performance that you are seeing?
- How does it evolve when using different threads? Why it scales very similarly than the polynomial evaluation?
 - Could you have a guess at the memory bandwidth of this machine?