



Modelling 2D steady-state heat equation

PTfS-CAM Project

Part II

PTfS 2020



Solve for u : $A u = f$

1. Use Conjugate Gradient (CG)
2. Use Preconditioned Conjugate Gradient (PCG) with symmetric Gauss-Seidel preconditioning



1. Calculate roofline predictions in [LUP/s] for CG and PCG solvers on 1 socket (10 cores) of Emmy. Calculate for two grid sizes : 1000×20000 and 1000×100000 . Does the performance change? Why?
2. Check whether you attain the roofline performance by running the code on 1 socket of Emmy. Timings are already included. Run the code using following command.

```
./perf num_grids_y num_grids_x
```
3. Measure the code balance in [bytes/LUP] of 'APPLY_STENCIL' kernel on 1 socket (10 threads) for both the grid dimension and comment whether it agrees with your model. You can use LIKWID for the measurement.
 - To switch on LIKWID measurement set the LIKWID flag to 'on', i.e.,

```
LIKWID=on CXX=icpc make
```
4. How does the code scale from 1 socket to 1 node (2 sockets, 20 cores) of Emmy? Does it scale perfectly? If not Why?

Chained roofline model



PCG example

$$r = b - \mathbf{A} x$$

$$\text{res_norm} = \langle r, r \rangle$$

$$z = \mathbf{P} r_0$$

$$\alpha_0 = \langle r, z \rangle$$

$$p = z$$

```
while( (iter<niter) && (res_norm > tol*tol) )
```

$$v = \mathbf{A} p$$

$$\lambda = \frac{\alpha_0}{\langle v, p \rangle}$$

$$x = x + \lambda p$$

$$r = r - \lambda v$$

$$z = \mathbf{P} r$$

$$\alpha_1 = \langle r, z \rangle$$

$$p = z + \frac{\alpha_1}{\alpha_0} p$$

$$\alpha_0 = \alpha_1$$

```
++ iter
```

Multiple kernels

Chained roofline model – a simple example



```

void scale(int n,...)
{
    for(int i=0; i<n; ++i)
        a[i] = c*b[i];
}

void add(int n,...)
{
    for(int i=0; i<n; ++i)
        a[i] = b[i] + d*c[i];
}

int main()
{
    ...
    scale(1e8,...);
    add(1e8,...);
    ...
    return 0;
}
    
```

$$I_{\text{scale}} = \frac{1 \text{ Flop}}{24 \text{ Byte}}, \quad P_{\text{scale}}^{\text{max}} = \frac{1}{2} * 4 * 2.2 * 10 = 44 \text{ GFlop/s}$$

ST SIMD f CORES

$$P_{\text{scale}} = \min(44, 40/24) = 1.67 \text{ GFlop/s}$$

$$I_{\text{add}} = \frac{2 \text{ Flop}}{32 \text{ Byte}}, \quad P_{\text{add}}^{\text{max}} = \frac{1}{2} * 2 * 4 * 2.2 * 10 = 88 \text{ GFlop/s}$$

$$P_{\text{add}} = \min(88, 40*2/32) = 2.5 \text{ GFlop/s}$$

} Is it $P_{\text{total}} = P_{\text{scale}} + P_{\text{add}}$? **NO**
 But $T_{\text{total}} = T_{\text{scale}} + T_{\text{add}}$

$$T_{\text{total}} = \frac{1e8*1}{P_{\text{scale}}} + \frac{1e8*2}{P_{\text{add}}} = 0.1398 \text{ s}$$

$$\Rightarrow P_{\text{total}} = \frac{1e8}{T_{\text{total}}} = 0.7 \text{ GIT/s}$$

Things to take care of



- Use EMMY (Ivy Bridge) for getting your performance results.
- Check if the measurements are reproducible. (Think of **pinning**, **scheduling**, and **clock frequency**).
- Remember pinning (`-C` with `likwid-perfctr`) when doing performance counter measurements. Use `-m` for using markers, instead of end-to-end measurement.
- When measuring parallel loops with `likwid-perfctr`, `LIKWID_MARKER_START` and `LIKWID_MARKER_STOP` should be called by all threads.
- Remember `likwid-perfctr` can incur some overheads, so for performance measurement run without `'LIKWID=on'`.
- Remember to arrive at final roofline model of (P)CG, you would need to stitch performance models of different kernel. It might be convenient to use time based model (see previous slide).



- Submit the code after doing tasks until **July 27, 2020** at PTfS Moodle (link will be available). Please compress and submit only a single (ZIP) file. Include code and the report in submission.
- The report should contain all the details necessary to reproduce your measurements.
- All submitted code should be compilable by just typing `make` (Makefile is already provided),
- Your code should pass all the tests.
- Both executables (`test` and `perf`) should run without segfaults.
- While submitting report **expected roofline performance** and the **measured performance**, use $\left[\frac{LUP}{s}\right]$ ($= \left[\frac{IT}{s}\right]$) as performance metric, see definition in `perf.cpp`.
- If there is any substantial deviation between these values, please explain plausible cause if any.



- In the exam you will be definitely asked questions based on this exercise.
- Happy Coding !!!