FFF2: Future-proof High Performance Numerical Simulation for CFD with FEASTFLOW (2)

Markus Geveler, Dirk Ribbrock

Inst. f. Applied Mathematics, TU Dortmund
(markus.geveler@math.tu-dortmund.de)

Principal Investigator

Prof. Dr. Stefan Turek
Inst. f. Applied Mathematics
TU Dortmund

Description of the DECI work

In our *hardware-oriented numerics* policy, we generally aim at the development of efficient, reliable and future-proof numerical schemes for the parallel solution of partial differential equations (PDEs) arising in industrial and scientific applications. Here, we are especially interested in *technical flows* which can be found in a wide variety of (multi-) physics problems. In our approach, both *numerical* and *hardware* efficiency are addressed simultaneously: Algorithms and whole solvers have to be tailored with respect to the target hardware in order to achieve a significant amount of the parallel peak performance. On the other hand, the sole concentration on hardware efficiency does not carry out the whole job (and in some cases may be counter-productive): Numerical efficiency plays a crucial role and itself includes multiple levels that can be optimised. Starting with the overall numerical and algorithmic approach required for the solution of a given domain specific problem (i.e. discretisation of the governing equations in time and space), stabilisation, linearisation of non-linear problems and finally the solution of the linear problems and smoothing therein, all these aspects together with the aforementioned levels of parallelism bear a large amount of interdependencies.

We have improved our recently very successful, parallel numerical software framework FEASTFLOW[1]. The improvements take both aspects of efficiency into account in order to make it ready for future HPC architectures: Novel numerics- and physics-components as well as software techniques for massively parallel (heterogeneous) compute resources have been employed that extend the applicability of the package to current and future real-world problems in the field of CFD. In the time when FFF was launched, the majority of the codes where under development or had undergone significant framework augmentations aiming at utilising the heterogeneous compute nodes of the CINECA PLX cluster. Hence the project's major focus had been framework development and testing, with some application production runs. The complementary PRACE-PA proposal for FFF-NGA has been solely intended to explore the capabilities of the XEON Phi accelerator in our hybrid solvers. This follow-up project, FFF2, was supposed to continue the basic research of FFF but with more focus on whole applications.

With the resources provided for FFF2 by the PLX cluster at CINECA, we decided to go for the optimisation of medium-scale solver runs employing multi-threading and GPGPU acceleration. This time, we concentrated on *free-surface flow simulation with pollutant transport in order to have a sort of multi-phase flow simulation*, based on the coupling of the Shallow Water Equations,

$$\frac{\partial h}{\partial t} + \frac{\partial (hu_j)}{\partial x_j} = 0 \quad \text{and} \quad \frac{\partial hu_i}{\partial t} + \frac{\partial (hu_i u_j)}{\partial x_j} + g\frac{\partial}{\partial x_i}(\frac{h^2}{2}) = S_i^{\text{bed}} + S_i^{\text{wind}}$$

and a convection-diffusion equation

$$\frac{\partial hc}{\partial t} + \frac{\partial (hu_j c)}{\partial x_j} = \frac{\partial}{\partial x_j}\left(Dh\frac{\partial c}{\partial x_j}\right) + S^{\text{poll}}$$

where

$$S_i^{\text{bed}} = -g\left(h\frac{\partial b}{\partial x_i} + n_b^2 h^{-\frac{1}{3}} u_i \sqrt{u_j u_j}\right) \quad ,$$

$$S_i^{\text{wind}} = (\rho_\alpha 10^{-3} \times (0.75 + 0.0067\sqrt{w_1^2 + w_2^2}))(w_1\sqrt{w_1^2 + w_2^2})$$

and

$$S^{\text{poll}} = -Khc + S_0 h \quad .$$

The coupled numerical schemes in use stem from the augmentation of earlier versions[2, 3], which are based on the Lattice-Boltzmann method and which are yet to be published. The resulting application is a common representative of a low-latency neighbour-to-neighbour communication-heavy stencil type code.

Results achieved by the DECI work

The development of the numerical components had already begun in the predecessor project FFF. In FFF2 we concentrated on optimisation of the kernels, datastructures and framework. Results of a typical simulation setup are displayed in Figure 1.

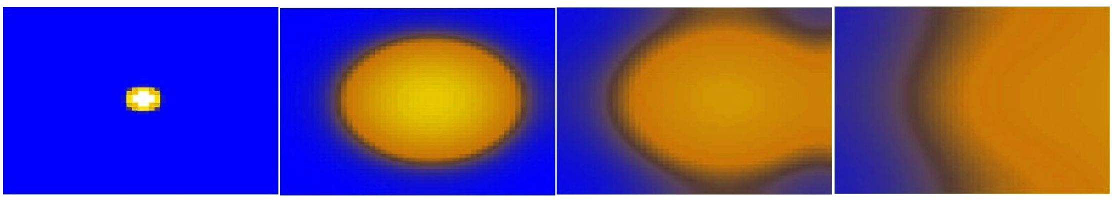

Figure 1: Pollutant distribution over time in a channel (oil in water). Constant uniform flow from the left, outlet at the right and high gradient gaussian initial concentration of the pollutant.

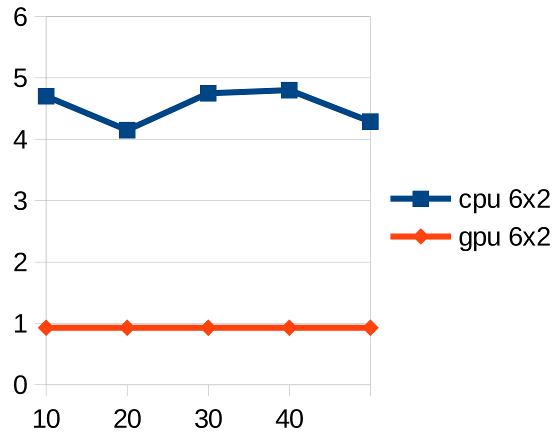

#procs	cpu	gpu
1	17.6	
2	8.9	
4	4.5	0.9
8	2.5	0.46
16	1.2	0.24
32		0.25

Figure 2: Weak (left) and strong (right) scaling survey of the prototype application using 2 MPI processes per node with 6 OpenMP threads and one GPU each, resulting in 2x6 threads or 2 GPUs per node respectively. In the weak case, 4.2x10^6 lattice sites are used per node. In the strong case, 33.5x10^6 unknowns are employed initially.

Particularly, in the course of the project, we were able to develop, optimise and employ several new features:

→ A preprocessor-based auto-loop-unrolling stage providing SIMD-width-oblivious implementation of all kernels.

→ Overlap of GPU up-/ downloads with computations using CUDA streams.

→ Overall use of CUDA-aware MPI if possible.

→ Parallel implementation of the domain partitioning and setup-/ preprocessing stages as opposed to file-based setup. Here, one-sided MPI communication was used to distribute the data to the nodes. This removes the bottleneck of a single node's memory size when partitioning large grids.

→ Overall use of device-to-device transfer between GPUs if possible.

Results from a scalability survey on the CINECA PLX cluster can be found in Figure 2. As can be seen, scalability is sustained despite the high level of optimisation. In the strong scaling of the GPU runs, at a total number of 32 processes, the problem becomes too small to generate additional speedups, as expected.

All described techniques are permanently employed into the next generation of FEASTFLOW.

Benefits of using the PRACE infrastructure

PRACE platforms employed

CINECA PLX

Title and goals of overarching projects

This work was supported (in part) by the German Research Foundation (DFG) through the Priority Programme 1648 `Software for Exascale Computing' (SPPEXA) which aims at tackling the various challenges of large scale scientific computing as well as programme TU 102/50-1. This work has on the other hand induced benefits for these projects.

References

[1] FEATFLOW project homepage, http://www.featflow.de

[2] Geveler, M.; Ribbrock, D.; Mallach, S.; Göddeke, D.; Turek, S., *A Simulation Suite for Lattice-Boltzmann based Real-Time-CFD Applications Exploiting Multi-Level Parallelism on modern Multi- and Many-Core Architectures*, Journal of Computational Science, 2, 113-123, doi 10.1016/j.jocs.2011.01.008, 2011

[3] Göddeke, D.; Komatitsch, D.; Geveler, M.; Ribbrock, D.; Rajovic, N.; Puzovic, N.; Ramirez, A., *Energy efficiency vs. performance of the numerical solution of PDEs: an application study on a low-power ARM-based cluster*, J. Comput. Phys., 237, 132-150, DOI 10.1016/j.jcp.2012.11.031, 2013