Date 09/06/2023
Contact person Dr. Matthias Möller
Our reference
Telephone +31 (0)15 27 89755
E-mail m.moller@tudelft.nl
Subject Referee's report on the PhD thesis Jakub Klinkovský

Delft University of Technology

Faculty EEMCS
Delft Institute for Applied Mathematics

Visiting address
Mekelweg 4
2628 CD Delft

Postal address
PO box 5031
2600 GA Delft
The Netherlands

http://math.ewi.tudelft.nl

To
Mrs. Monika Zabranska
CVUT – FJFI
Brehova 7
115 19 Praha 1
Czech Republic

Dear Prof. doc. Ing. Cuba,

I am sending you my evaluation of the PhD thesis by Ing. Jakub Klinkovský entitled

Data Structures and Parallel Algorithms for Numerical Solvers in Computational Fluid Dynamics

I have carefully studied the manuscript and ***recommend accepting it for presentation and defense***. In what follows I will substantiate my recommendation in more detail.

The dissertation consists of 7 well-written chapters followed by conclusions and 7 appendices. The dissertation's content can be structured at coarse level into:

i) Parallel programming frameworks, data structures for unstructured data, and linear solvers (Ch. 1-3);
ii) mixed-hybrid finite element method (Ch. 4);
iii) lattice Boltzmann method (Ch. 5); and
iv) the combination of MHFEM and LBM and its application to vapor transport in turbulent air flow (Ch. 6-7).

The first three chapters are mainly focusing on software and computational aspects with special attention to the implementation on CPUs and GPUs.

In particular, **Chapter 1** provides very valuable information on how a simple but frequently reoccurring compute kernel – axpy – can be implemented in today's parallel programing frameworks. While the provided source-code snippets give a good impression of *how* to implement the compute kernel, Chapter 1 would have benefitted from a performance comparison of the different implementations on CPUs and GPUs. In its current form, Chapter 1 only helps the reader to appreciate that some frameworks – like OpenMP/OpenACC – require only very minimal code changes whereas others – like OpenCL – require extra boilerplate code. That said, the presented material is definitely up to date and covers the relevant parallel programming frameworks that are in use today.

Chapter 2 continues with the development of efficient data structures for handling conforming unstructured meshes on distributed compute infrastructures, in particular, multiple GPUs and their implementation in the open-source TNL library that is being developed by the PhD candidate and his supervisor. From the mathematical point of view, this is not 'rocket science' but, at the same time, it is one of the most crucial building blocks for achieving high computational performance in subsequent simulations. What makes this chapter particularly interesting from the scientific point of view is the description of why certain design considerations were made (cf. Section 2.2.4), which can be of help for the reader in making their own design choices. Next to the description of the data structures, Chapter 2 contains a detailed benchmarking of the in-house implementation in TNL and an existing reference implementation in MOAB. This description is very detailed and written in a way that allows the reader to reproduce the results and/or extend to benchmark to other libraries. I consider this very good scientific practice and a very valuable contribution of the dissertation.

Chapter 3 falls somewhat short as it just describes the linear solvers available in TNL and competitors' libraries but does not provide any stand-alone performance benchmarking. That said, the solvers' performance is tested implicitly in subsequent chapters in the context of concrete numerical methods.

The remaining chapters of the dissertation represent the `mathematical core' of the thesis work. **Chapter 4** describes the mixed-hybrid finite element method (MHFEM) that is a continuation of the candidate's master project's work. The main difference is the reimplementation of the method in the TNL library and the utilization of the novel distributed mesh data structures that allows the candidate to perform large-scale computations on up to 14 MPI nodes with 336 cores and 4 GPUs, respectively. As in Chapter 2, the candidate provides an excellent description of the benchmarking methodology that lends itself for reproducing the results presented in the dissertation. The interpretation of the benchmark results is less straightforward. The results from the strong scaling analysis shed some light on how the different implementations (OpenMP vs. MPI) and linear solvers (BiCGStab solver with Jacobi preconditioner in TNL vs. BiCGStab solver with BoomerAMG preconditioner in Hypre) scale with an increase of the compute resources. However, it remains somewhat unclear what the candidate wants to achieve with this comparison. Looking at the compute times (CT) the Hypre-based solver stack (BiCGStab + BoomerAMG) is favorable over TNL's BiCGStab solver - Jacobi preconditioner approach which is no surprise since Jacobi is a quite weak preconditioner. By comparing two different solver stacks *and* solving strategies it is not easy to pinpoint if the performance differences are mainly caused by `better mathematics' (i.e. the better preconditioner that leads to fewer iterations) or `smarter implementation' (i.e. the faster allocation of memory, etc.). That said, it is still impressive to see that the computational efficiency (Eff) of the TNL stack is mostly en par with Hypre (cf. Table 4.5) and even exceeds it (cf. Table 4.4) in the MPI case.

Chapter 5 is the second `mathematical' chapter of the dissertation and it presents the results for the lattice Boltzmann method. As in previous chapters, the description of the benchmark methodology is excellent, and the computational efficiency demonstrated in the strong and weak scaling tests are convincing (cf. Table 5.3 for the strong scaling results and Tables 5.4-5.5 for the weak scaling results). What I particularly appreciate is the PhD candidate's scientific integrity and `honesty' in this but likewise in all other chapters. Jakub clearly states which results were produced by other team members and also reports, in the spirit of reproducibility, shortcomings and non-functional features of the hardware, e.g., the non-utilization of NVIDIA's GPUDirect technology on the Karolina supercomputer. Such details are often not reported in scientific literature but make it difficult to reproduce certain results.

Chapters 6 and 7 are the `mathematical highlights' of the dissertation with a coupled LBM-MHFEM approach applied to vapor transport in turbulent air flow. While in previous chapters, the focus was mainly on the performance of the computational building blocks, these chapters focus on the algorithmic/ mathematical performance of the developed hybrid model. As it becomes clear from Figures 7.2 and 7.8-7.13, the developed mathematical model is able to predict solutions that are relatively close to experimental results at least for certain set ups. Independent of any computational performance aspects, the proposed method is an interesting approach for solving multi-physics problems and, as such, has scientific value. The remaining parts of the dissertation – conclusions and appendices – provide valuable extra information.

In summary, I come to the conclusion that the dissertation has various scientific contributions both at the algorithmic/computational level and the methodological level. I therefor recommend to accept it for presentation and defense.

Sincerely yours

(Dr. rer. nat. Matthias Möller)