The 8th European Congress on Computational Methods in Applied Sciences and Engineering
ECCOMAS Congress 2022
5-9 June 2022, Oslo, Norway

Evaluation of the performance portability layer of different linear solver packages with ALIEN, an open generic and extensible linear algebra framework

Cédric Chevalier[1], Stéphane De Chaisemartin[2], Jean-Marc Gratien[2], Thomas Guignon[2], Pascal Havé[3] and Xavier Tunc[2]

[1] CEA, DAM, DIFF, F-91297 Arpajon, France http://www.cea.fr
[2] IFP Energies nouvelles, 1-4 avenue de Bois-Préau, 92852 Rueil-Malmaison, France
http://www.ifpenergiesnouvelles.fr
[3] HAVENEER, https://www.haveneer.com

Applications to solve large and complex partial derivative equation systems often rely nowadays on frameworks like Arcane [1], Dune [3], Feel++. Linear solver packages like PETSc or Trilinos are used to manage linear systems and to have access to a wide range of algorithms. With the evolution of High-Performance Computing, the variety of the hardware features available in new architectures has considerably increased. ARM processors, AMD, Intel and Nvidia GP-GPUs, TPU and FPGA devices are now common features. To handle the induced complexity, different strategies are adopted in each linear solver frameworks. Some of them consist in introducing a new layer that provides abstractions to manage the performance portability and to enable several parallel programming models.

In this paper, we evaluate the performance of linear solver packages that rely on tools like SYCL[6], Kokkos [4] or HARTS [5] to handle runtime systems like OpenMP, TBB, CUDA,... A simulator to solve advection-diffusion problems has been developed with ALIEN [2], a C++ framework that provides a high level and unified API to handle large distributed matrices and vectors. We have benchmarked different solver algorithms. We have evaluated the efficiency of their implementations, their capability to perform on different architectures with for instance large number of cores, GP-GPU accelerators or processors with large SIMD instructions.

REFERENCES

[1] Arcane framework : https://github.com/arcaneframework/arcane.git

[2] Alien framework : https://github.com/arcaneframework/alien.git

[3] Bastian, Peter, Heimann, Felix, et Marnach, Sven. Generic implementation of finite element methods in the distributed and unified numerics environment (DUNE). Kybernetika, 2010, vol. 46, no 2, p. 294-315.

[4] Edwards, H. C., Trott, C. R., Kokkos: Enabling performance portability across manycore architectures.In 2013 Extreme Scaling Workshop (xsw 2013) (pp. 18-24). IEEE.

[5] Gratien, J. M., An abstract object oriented runtime system for heterogeneous parallel architecture. In 2013 IEEE International Symposium on Parallel and Distributed Processing, Workshops and Phd Forum (pp. 1203-1212). IEEE.

[6] Alpay, A., Heuveline, V., SYCL beyond OpenCL: The architecture, current state and future direction of hipSYCL. In Proceedings of the International Workshop on OpenCL