Preparing a Fortran legacy code for the upcoming exascale architectures

Joeffrey Legaux[1] and Gabriel Staffelbach[2]

[1] CERFACS, Toulouse - France, joeffrey.legaux@cerfacs.fr
[2] CERFACS, Toulouse - France, gabriel.staffelbach@cerfacs.fr

Keywords: *Computational Fluid Dynamics, High Performance Computing, GPU, ARM*

AVBP is a parallel CFD code that solves the three-dimensional compressible Navier-Stokes on unstructured and hybrid grids. Its highlight is the prediction of unsteady reacting flows in combustor configurations based on the Large Eddy Simulation (LES) approach.

AVBP is a cutting-edge software when it comes to distributed memory X86 CPUs, efficiently scaling up to 200.000's of cores. However, other types of architectures such as ARM processors and accelerators are gaining popularity and are expected to play a significant role in the exascale era. Extending AVBP to support such architectures will soon become mandatory if we want to be able to efficiently use all high-end systems in the future.

AVBP is both a legacy Fortran code and a very active research project ; this implies a very large and quickly evolving code base, whose most developers have little to no expertise in the HPC domain per se. Porting it to new architectures thus cannot occur at the expense of its accessibility and maintainability, that is why we need to evaluate if new architectures can be efficiently exploited while still relying on the current CPU code base.

We first explore the usage of GPU accelerators through OpenACC directives. We tried various approaches to conduct this port, ultimately settling for a data-centric approach. We hoped the directive-based implementation would allow us to keep a single common code base for both CPU and GPU usages ; however keeping a completely unified code base could not provide ideal performance on both architectures and we had to introduce specific implementations for some parts. The OpenACC implementation of AVBP currently covers the majority of its use cases and its performance has been assessed on various HPC platforms including NVIDIA's V100, A100 and A30 GPUs.

Another architecture we explore is the ARM processors. Motivated by new and exciting possibilities such as Graviton (AWS) and Fujitsu A64FX, supporting these systems has become strategic since the European Processor Intiative interest on these architectures and the development of their own flavor. Porting to these systems has proven less challenging than GPUs since it depends more on compiler support, although leveraging performance after the initial port has proven more difficult and requires more extensive work. In this section we will highlight the porting difficulties and the intricacies of choosing the best compiler/coding scenarios to yield the best performance until these work enters the production phase.