


COSMO Physics Modules Are Accelerated as a Part of HP2C Initiative Using OpenACC



Introduction

COSMO is a local-area atmospheric model used for numerical weather prediction. It is maintained by the Consortium for Small-scale Modeling (a.k.a., COSMO), and used by seven national weather services, including the Deutsche Wetterdienst (DWD) and MeteoSwiss. Scientists at more than seventy universities and research institutes also use COSMO for climate research. COSMO contains about 300K lines of Fortran 90 code.

CHALLENGE

Increase the compute performance of COSMO in a two cabinet footprint by 40X while maintaining portability to X86 clusters and code readability and accessibility among weather and climate researchers

SOLUTION

Used OpenACC directives for physical parameterizations – where code sharing was paramount – and created a C++/CUDA-based DSL — platform-portable and easily integrated with OpenACC — for the dynamical core.

RESULTS

MeteoSwiss deployed a two-cabinet CS-Storm clustered supercomputer with 192 NVIDIA Tesla K80 GPUs and achieved a 40X improvement in compute performance and 3X greater energy efficiency, while maintaining portability to x86 clusters and accessibility and code sharing in the scientific community.

Challenges

Starting in 2010, developers from MeteoSwiss, ETH Zurich and CSCS re-engineered the COSMO model, as part of the Swiss HP2C initiative. The goals of the project were to:

- Increase the resolution of the operational weather simulations by 2X for improved daily and five-day forecasting, which mandated a 40X increase in compute performance.
- Adapt the COSMO model to run on energy-efficient accelerated systems, while maintaining portability to clusters of multi-core x86 CPUs.
- Maintain the readability and accessibility of the existing sources, for continued sharing of methods and modules across the scientific weather and climate community.

A numerical weather prediction model has three main parts: the dynamical core, the physical parameterizations and data assimilation. The dynamical core of COSMO (40K lines of code/60% of runtime) solves the equations of fluid motion and thermodynamics. Physical parameterizations (130K LOC/25% of runtime) model turbulence, radiative transfer and cloud microphysics. (Assimilation of weather data and radar information accounts for the remaining 15% of runtime.)

Accelerating COSMO: Physics and Dynamical Core

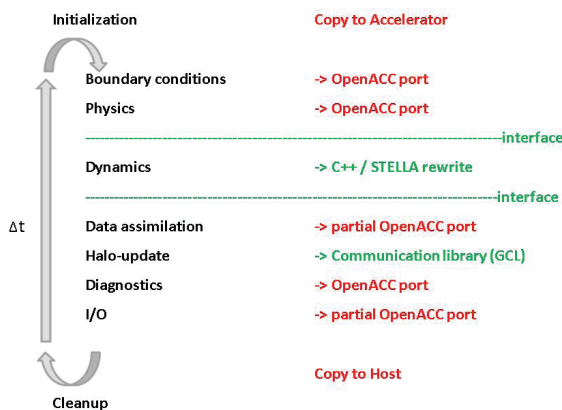
To speed-up the physical parameterizations, the HP2C team used OpenACC directives. OpenACC developers insert directives into their Fortran, C or C++ source code that help a OpenACC compiler identify and parallelize compute-intensive kernels for execution on accelerators. OpenACC directives look like comments to ordinary compilers.

CASE STUDY — COSMO DYNAMICS AND PHYSICS MODULES ARE ACCELERATED AS A PART OF HP2C INITIATIVE USING OPENACC

The physical parameterizations were characterized by low compute-intensity: frequent data transfers between host and accelerator memory were time-consuming. COSMO developers copied all the data onto the accelerator, and kept it resident using OpenACC `data` directives and the `present` clause. HP2C developers fused some imperfectly nested loops to avoid launching additional kernels (Figure 1).

The dynamical core uses the data already on the accelerator in the middle of each time step. The HP2C team aggressively rewrote the dynamical core – which had only a handful of developers – creating the STELLA DSL to implement the different stencil motifs. Cross-platform portability was ensured by creating both C++/x86 and CUDA back-ends. HP2C developers used the OpenACC host data directive to interoperate with the C++/CUDA dynamical core.

Figure 1: COSMO Workflow



Integration and Production

The GPU-accelerated COSMO 4.19 hybrid modules were integrated into the production COSMO 5.0 trunk in 2014. In April, 2016, MeteoSwiss became the first major national weather service to deploy GPU accelerated supercomputing for daily weather forecasting. MeteoSwiss needed the 40X performance improvement — shown in Figure 2 — in the footprint of their previous two-cabinet Cray XE6. MeteoSwiss deployed a Cray CS-Storm cluster supercomputer, Piz Kesch, in two cabinets, each with 96 NVIDIA Tesla K80 GPUs and 24 Intel Haswell CPUs. Piz Kesch is about three times as energy efficient and two times as fast as a contemporary, conventional all-CPU system.

Dr. Oliver Fuhrer — from MeteoSwiss — says OpenACC brought three important benefits:

- OpenACC made it practical to develop for GPU-based hardware while retaining a single source for almost all the COSMO physics code.
- OpenACC, a CUDA-based DSL, and GPUs can accelerate simulations while reducing the electric power required of the computer system executing the simulation.
- OpenACC makes it possible for domain scientists to continue normal scientific developments while creating a new GPU accelerated version of the COSMO model.

Figure 2: Breakthrough Advanced in Swiss Weather Forecasting

