

Numeca Taps OpenACC to Accelerate Commercial CFD Application without Rewriting Code



OpenACC
Directives for Accelerators

CHALLENGE

To meet growing customer demand for GPUs, Numeca wanted to accelerate the Fine/Turbo application on GPUs without adapting legacy code. This would provide faster time-to-solution for customers.

SOLUTION


OpenACC and the PGI Accelerator compiler enabled developers to restructure and accelerate computationally intensive routines, while maintaining portability and eliminating the need to rewrite code.

RESULTS

- With OpenACC delivering 10X or higher speed-ups on key routines, typical cases showed a full-application global speedup of up to 2X on the Oak Ridge Titan supercomputer.
- Total time spent on optimizing various routines with OpenACC was about five person-months.
- Customers can achieve more accurate results with the accelerated code in half the time, enabling much faster design cycles

Project Overview

Numeca is a Belgian-based software company providing high-fidelity computational fluid dynamic (CFD) tools (www.numeca.com). With roughly 120 employees and 70 software developers, the company's Fine/Turbo software product—a structured, multi-block, multi-grid CFD solver targeting the turbo machinery industry—has more than 2,000 users in companies such as Honeywell and multiple supercomputing centers.



Working with Oak Ridge National Laboratory, the software development team wanted to accelerate the code on the Titan supercomputer while maintaining portability.

“The target platform was Titan at Oak Ridge Leadership Computing Facility, but we were also interested in doing a general port which could be sold to customers as a commercial product,” said David Gutzwiller of Numeca.

Challenge

According to Gutzwiller, there were numerous challenges in accelerating Fine/Turbo, including:

- **Development time:** Only a few person-months were allocated for adding GPU acceleration into the code, which required a portable approach with minimal developer effort.
- **Code maintenance:** Fine/Turbo is in constant development by a team of engineers. The GPU acceleration should not interfere with the ability of other developers to modify and commit changes or bug fixes. Also, there should be no duplicate sources for easy maintenance.
- **Code portability:** Many Fine/Turbo customers still use traditional homogenous CPU systems. The GPU acceleration work should be portable on those systems.

“Writing and maintaining a code base of this size is a challenge in itself, and requires a huge outlay of effort to make large-scale changes,” he said. “We were on a tight schedule, and we needed high accuracy for the end product, so there was little room for error.”

“OpenACC enabled us to target routines for GPU acceleration without rewriting code, allowing us to maintain portability on a code that is 20-year old.”

— David Gutzwiller, Numeca

Solution


Leveraging OpenACC to repurpose the code, Gutzwiller's team followed the host-accelerator paradigm. Thirty of the most computationally expensive routines were targeted for acceleration. Host processes on the CPU manage the computation, offloading the high-value routines to the GPU for improved execution time.

"OpenACC enabled us to target and prepare routines for GPU acceleration without rewriting them which helps to minimize potential errors," said Gutzwiller. OpenACC also enabled the team to prepare the code to run on many different types of supercomputers, enabling portability across all customer environments.

Results

With OpenACC, key routines often saw 10x speed-up or more compared to highly tuned CPU implementation. Typical configurations have shown a global speed-up of about 2X, when 16 CPUs are running on the OLCF Titan system (16 core AMD CPU + 1 NVIDIA Tesla K20 per node), depending on the model configuration and mesh size. The best observed speed-up on Titan is 2.5X with a very clean, well-balanced model. Because the 30 routines make up roughly 70 percent of the run time, 1.75-2.5x speed-up surpassed expectations. Total time spent on accelerating Fine/Turbo with OpenACC was about five person-months.

GPU acceleration with OpenACC resulted in a large improvement in time-to-solution, enabling a much faster engineering design cycle, which, according to Gutzwiller, is very appealing to both researchers and industrial customers. The accelerated code will enable customers to cut costs significantly, while speeding time to market with new products. "On a supercomputer, time is the biggest cost," he said. "With the acceleration, users spend half the hours to get the same results."



Gutzwiller believes the majority of Numeca's customers will leverage the accelerated code moving forward. He plans to expand instrumentation with OpenACC further into other Fine/Turbo modules and Numeca tools. "Most of our customers have GPUs," he said. "Once we get to all our customers, there's no telling how many will use it and reap the benefits."

Gutzwiller will present his findings at a workshop at Supercomputing 2015 with his paper, "Acceleration of the FINE/Turbo CFD Solver in a Heterogeneous Environment with OpenACC Directives." [Find out more](#)

