

CAPS offers services to help you build optimized applications running on parallel high performance systems. Ranging from training to complete application porting, we want to give you all the expertise your problem might require.

We have a large experience in using CUDA™ to program the NVIDIA® Tesla™ accelerators in a portable and interoperable way.



CUDA™ Training

We usually operate in your office for a group of 10 persons mixing theoretical courses with practical exercises:

- **Basic Programming**

All the basic concepts that allow you to start playing with CUDA™.

- **Advanced User Programming**

For advanced users, this training focuses on optimizing CUDA™ kernels running in parallel with the main application.

The practical exercises are based on your domain-specific algorithms.

- **Parallel and Heterogeneous Programming**

In this course, all different levels of parallelism are studied: MPI, OpenMP, CUDA™.

The goal of the course is to give an overview of how to parallelize applications in a coherent and efficient way.

CUDA™ Engineering

Based on a study of your application, we first evaluate the computations that will benefit from a NVIDIA® Tesla™ acceleration. We then provide you with the CUDA™ kernels and their integration glue in your application so as to minimize the data transfers and maximize the performance.

CUDA™ Expertise

Ask us for any issue you might have in using CUDA™:

- Study of an appropriate hybrid machine configuration
- CUDA™ kernel tuning
- Application mapping strategy on a parallel hybrid cluster

