Internship in data science for scientific computing: Accelerating advanced stochastic methods on modern computer architectures.

Cerfacs, Toulouse.

Position: M2 or engineer internship
Background: M1 (or equivalent) in computer science or applied maths
Starting date: March 2020 or later
Duration: 6 months
Salary: 650€ net per month
Hosting institution: Cerfacs (Toulouse), https://cerfacs.fr
Contact persons: Paul Mycek (Cerfacs) and Marc Baboulin (University Paris-Sud)
Keywords: Data science, uncertainty quantification, high performance computing, GPU programming, tensor cores, mixed precision arithmetics, multilevel Monte Carlo.
How to apply: Send a CV and a cover letter to mycek@cerfacs.fr and baboulin@lri.fr.

Description

Multilevel Monte Carlo (MLMC) sampling is an approach popularized in 2008 by Giles for the solution of stochastic differential equations. It has since been adapted to accelerate **uncertainty quantification** (UQ) sampling methods for **partial differential equations** (PDEs), with applications in various areas of engineering and applied mathematics for industrial problems. The underlying idea of MLMC is to take advantage of different levels of numerical discretization in such a way that many (cheap) evaluations are performed on the coarsest levels while fewer (expensive) computations are required on the finest levels, resulting in a reduced computational cost. In terms of error, many coarse evaluations help reduce the sampling error while fewer fine evaluations help reduce the discretization error. MLMC methods have a solid mathematical foundation and have demonstrated their superiority over single-level Monte Carlo sampling methods on a wide range of applications.

Modern computer architectures, especially upcoming **exascale parallel systems**, provide exciting opportunities to design efficient MLMC methods that make the most of **mixed precision floating-point arithmetics** offered by hardware accelerators such as **graphical processing units** (GPU) and **tensor cores**. This is indeed very well aligned with the MLMC core idea of balancing different sources of errors. The proposed internship project aims at developing a prototype MLMC algorithm exploiting these concepts.

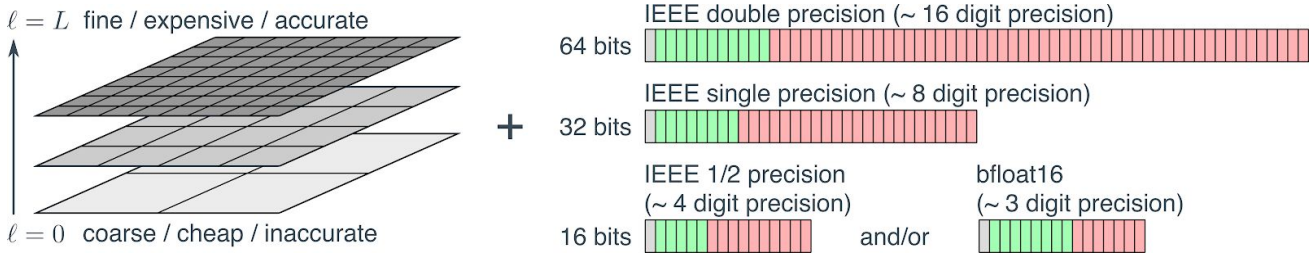

Required skills

- Having completed one parallel programming or HPC course
- Experience in C or C++ programming

Desired skills

- Experience with GPUs
- Experience in scientific computing