Resilient PDE solving approaches for exascale computing

Paul Mycek

CERFACS

Toulouse, France

Abstract

Resilience represents one of the many challenges to be anticipated and addressed before viable exascale systems can emerge. In this context, I will present an approach to tackle resilience for the solving of deterministic and stochastic partial differential equations (PDEs).

The approach is based on a domain decomposition framework, together with a sampling and optimization approach. In particular, the occurrence of silent data corruption (SDC), modeled by random bit flips in the 64-bits representation of double precision variables, is overcome by means of robust regression techniques involving l1-minimization. In the case of an elliptic PDE with random coefficients, the stochastic quantities are approximated by a polynomial chaos (PC) surrogate and we then resort to a hybrid approach involving both sampling and Galerkin projection.

Finally, I will present discrete a priori bounds on the solution of elliptic PDEs, used to enhance the resilient capabilities of the approach. This a priori knowledge is used to check the admissibility of local PDE solutions and to filter out corrupted data, thus making the subsequent regression more robust.