CEMRACS 2016: project ENLAK
Enlarged Krylov solvers applied to industrial problems

Leader: Laura Grigori, Inria Paris, Alpines

Participants:

- Hussam Al Daas, Inria Alpines/Total
- Laura Grigori, Inria Alpines
- Olivier Tissot, Inria Alpines
- Pascal Henon, Total
- Philippe Ricoux, Total
- Yvan Fournier, EDF

Funding: Total, Inria Alpines, NLAFET FETHPC H2020 project, EDF

Abstract: One of the main challenges in high performance computing is the exponentially growing gap between the time required to perform floating point operations and the time required to move data between processors. The goal of this project is to analyze Krylov subspace methods that reduce communication in the context of two industrial applications from Total and EDF.

Presentation of the project Several works show that there is an exponentially increasing gap between the time required for solving floating point operations and the time required to transfer data, either between different levels of the memory hierarchy or between different processors. This gap is already seen in the current, highly optimised applications, from EDF and Total. Typically, when a linear solver based on an iterative method is run on more than several thousands of cores, the communication cost grows very quickly, eventually dominating the runtime of the solver.

In a quest to address this problem, Alpines Inria group focuses on designing communication avoiding methods, that is methods that reduce the communication to a minimum. Krylov subspace methods are largely used for solving linear systems of equations. In [1] we introduce enlarged Krylov methods, that enrich the basis by splitting the residual among t subdomains, thus obtaining a set of t vectors $T(r_0)$. The enlarged Krylov subspace is $K_t(A; r0) = span[T(r_0), AT(r_0), A^2T(r_0)\ldots]$, where the $T(r_0)$ is the set of t vectors. In the following iteration, instead of multiplying A with a vector, the new method computes the product of A with a set of t vectors. The communication remains still between neighboring processors as in classic Krylov methods, while the matrix-set of vectors operation has a higher arithmetic intensity than the matrix-vector operation. And as a result, t new vectors are added to the search space instead of a single one. The enlarged Krylov methods converge significantly faster than the classic Krylov methods.

Objectives In this project we will consider two industrial applications. The first application from Total concerns the reservoir simulation, which is a crucial tool for the development of oil and gas fields. It aims at predicting how much hydrocarbons can be recovered, optimizing the field development (well position, pressure of injection etc..) and testing new Enhanced Oil Recovery mechanisms. Having access to efficient and scalable sparse linear solvers is key for the reservoir simulator.

The second application from EDF concerns the CFD solver implemented in code Saturne, solving the Navier-Stokes equations for incompressible flows. An important part of the time in this code is spent in the linear solvers. This code is used for a wide range of applications, related for example to nuclear engineering as well as renewables.

The objective of this project is to study enlarged Krylov methods in the context of the two industrial applications. Starting from the work in [1] and the recent work of Hussam Al Daas and Olivier Tissot, our goal is to evaluate the efficiency of enlarged Krylov methods for these two considered industrial applications and extend them accordingly.

References

[1] L. Grigori, S. Moufawad, and F. Nataf. Enlarged krylov subspace conjugate gradient methods for reducing communication. *SIAM Journal on Matrix Analysis and Applications*, 37(2):744–773, 2016.