11th World Congress on Computational Mechanics (WCCM XI)
5th European Conference on Computational Mechanics (ECCM V)
6th European Conference on Computational Fluid Dynamics (ECFD VI)
July 20–25, 2014, Barcelona, Spain

PARALLELISATION STRATEGIES FOR THE PROPER GENERALIZED DECOMPOSITION ON MASSIVELY PARALLEL ARCHITECTURES (GPU).

D. Borzacchiello[1], A. Leygue[1], F. Bordeu[1] and F. Chinesta[1]

[1] Institut de Recherche en Génie Civil et Mécanique, UMR 6183, École Centrale de Nantes, 1 rue de la Noë BP 92101 44321 Nantes cedex 3

Key words: *Proper Generalized Decomposition, Reduced-order modelling, High-dimensional problems, GPU Parallel Computing.*

The Proper Generalized Decomposition (PGD) is an order reduction model originally introduced in the field of computational rheology by Ammar et al. [1, 2] to solve the Fokker-Planck equation describing the relaxation and orientation dynamics of polymer chains in a complex flow. The unparalleled performances offered by this technique for the resolution of high-dimensional problems are due to its intrinsic linear scaling with the dimensionality of the problem, since the PGD expresses the solution of a given equation in a N-dimensional space as a linear combination of appropriate terms, each being the tensor-product of lower dimensional functions. This feature has allowed the extension of the method to a variety of others fields ranging from quantum mechanics to the stress analysis of composite structures. The effective applicability of the PGD to many practical engineering problems has also encouraged the investigation of its theoretical background which is now well established and extensively documented by the latest works in the review literature [3, 4].

However, some inefficiencies that may affect PGD stem from problems whose solutions can be found in a separated representation, while the equations source terms or coefficients are not representable in a separated form. Even though these problems are by definition good candidates for PGD, in this case most of the solution time is spent to evaluate multidimensional integrals. The computational cost of this task can potentially overcome the one associated to resolution of the problem itself, reducing in this way the overall efficiency. A second critical point is associated to the fact that the fixed point algorithm used to solve the non-linear problem arising from the PGD formulation is intrinsically sequential and becomes suboptimal as the dimensionality of the problem increases, since it does not fully exploit the hardware computing capabilities.

Nevertheless, the level of maturity that has been reached in reduced-order modelling and the advent of General Purpose Graphic Processing Units (GPU) in the world of parallel

computing hardware, allows to rethink the implementation of the PGD and to experiment strategies for algorithmic optimization based on parallelisation. Indeed, GPUs allow an unprecedented level of parallelism on shared memory devices that can be hosted even on ordinary laptops with relatively low energy consumptions. The ultimate challenge is therefore to efficiently couple order reduction modelling with massive parallelism on energy efficient architectures, such as GPUs, to further reduce the cost of numerical simulation in scientific computing. The goal of the present communication is to explore the possible modifications of the classical algorithm of PGD in order to develop a massively parallel version.

Our research is specifically aimed to track and remove the bottlenecks present in the current implementation of PGD. In the first place, we seek to improve the performances of the fixed point algorithm by proposing a modified version consisting in the asynchronous update of all the lower dimensional functions at once. In this way the convergence rate is slightly lower, but on the other hand all the updates can be executed in parallel. The second proposed improvement, for the case of non-separable coefficients, is to perform the multidimensional integration using GPU acceleration, since reduction-type operations can be significantly sped up in this way. The scenario of possible applications of the parallel PGD includes the numerical computation of multi-parametric solutions for problems depending on a considerable amount of parameters. Indeed, this can be the case of multivariate shape optimization, a problem often encountered in aerodynamics or structural design applications, for which the evaluation of the objective function may require the resolution of a set of partial differential equations depending on the parameters to optimize. Since the coefficients of the governing equations are most likely non separable functions of these parameters, parallel PGD becomes the only viable option to circumvent both the high-dimensionality and the non-separability issues. In this way the solution can be pre-computed in an off-line step prior to the actual optimization algorithm.

REFERENCES

[1] A. Ammar, B. Mokdad, F. Chinesta, and R. Keunings. A new family of solvers for some classes of multidimensional partial differential equations encountered in kinetic theory modeling of complex fluids. *J. Non-Newtonian Fluid Mech.*, 139(3):153–176, 2006.

[2] A. Ammar, B. Mokdad, F. Chinesta, and R. Keunings. A new family of solvers for some classes of multidimensional partial differential equations encountered in kinetic theory modelling of complex fluids. *J. Non-Newtonian Fluid Mech.*, 144(2-3):98–121, 2007.

[3] F Chinesta, A Ammar, A Leygue, and R Keunings. Journal of Non-Newtonian Fluid Mechanics An overview of the proper generalized decomposition with applications in computational rheology. *J. Non-Newtonian Fluid Mech.*, 166(11):578–592, 2011.

[4] F. Chinesta, R. Keunings, and A. Leygue. *The Proper Generalized Decomposition for Advanced Numerical Simulations: A Primer*. Springer Publishing Company, Incorporated, 2013.