Accepted Manuscript

GPU accelerated computational homogenization based on a variational approach in a reduced basis framework

Felix Fritzen, Max Hodapp, Matthias Leuschner

PII: S0045-7825(14)00151-0
DOI: http://dx.doi.org/10.1016/j.cma.2014.05.006
Reference: CMA 10240

To appear in: *Comput. Methods Appl. Mech. Engrg.*

Received date: 27 September 2013
Revised date: 9 May 2014
Accepted date: 14 May 2014

Please cite this article as: F. Fritzen, M. Hodapp, M. Leuschner, GPU accelerated computational homogenization based on a variational approach in a reduced basis framework, *Comput. Methods Appl. Mech. Engrg.* (2014), http://dx.doi.org/10.1016/j.cma.2014.05.006

GPU accelerated computational homogenization based on a variational approach in a reduced basis framework

Felix Fritzen*, Max Hodapp and Matthias Leuschner

Young Investigator Group Computer Aided Material Modeling
Chair for Continuum Mechanics, Institute of Engineering Mechanics
Karlsruhe Institute of Technology (KIT), Kaiserstr. 10, D-76131 Karlsruhe, Germany
*felix.fritzen@kit.edu (corresponding author)

ABSTRACT. Computational multiscale methods such as the FE^2 technique (Feyel, 1999) come along with large demands in both CPU time and memory. In order to significantly reduce the computational cost of multiscale methods the authors recently proposed a hybrid computational homogenization method for viscoplastic materials using a reduced basis approach in a mixed variational formulation (Fritzen and Leuschner, 2013). In the present contribution two extensions of the method are introduced: First, the previous proposal is extended by allowing for heterogeneous hardening variables instead of piecewise constant fields. This leads to an improved accuracy of the method. Second, a massively parallel GPU implementation of the algorithm using Nvidia's CUDA framework is presented. The GPU subroutines for the batched linear algebraic operations are integrated into a specialized library in order to facilitate its use. The impact of the heterogeneous hardening states on the accuracy and the performance gains obtained from the dedicated GPU implementation are illustrated by means of numerical examples. An overall speedup in the order of 10^4 with respect to a high performance finite element implementation is achieved while preserving good accuracy of the predicted nonlinear material response.

KEYWORDS: Nvidia CUDA; graphics processing unit (GPU); GPU accelerated batched BLAS; reduced basis model order reduction; Generalized Standard Material (GSM); mixed incremental variational approach;

1 Introduction

Modern engineering mechanics simulations are dedicated to the prediction of the mechanical behavior of advanced materials. New types of lightweight construction materials, e.g. fiber reinforced composites in aircraft industry, seem to be homogeneous from a macroscopic point of view although they display strong heterogeneities on smaller scales. These inhomogeneities on the microscale influence the effective macroscopic mechanical behavior. The determination of the effective material properties, based on evaluations on the microscale, is called homogenization. The underlying analytical theory is well-documented by a huge body of literature (see, e.g., the compendium of Nemat-Nasser and Hori, 1999). However, analytical methods have their limitations for increasing geometric complexity of the microstructure and when it comes to accurate predictions of nonlinear material behavior.

The fast-growing information technology sector and the attendant increasing computing capacities gave rise to the development of appropriate tools in the field of computer assisted mechanical simulations. Especially the well-known finite element method (FEM, Bathe, 1996) has been proven as a powerful tool in computer aided engineering. The use of a nested FEM approach for homogenization problems was introduced by Feyel (1999) in terms of the nested finite element method FE^p. Thereby, a discretized microstructure is assigned to every integration point on the macroscale and solved numerically using an individual FEM simulation. However, this method results in massive computational effort and tremendous memory requirements, even for two-dimensional problems.

The unacceptable computational cost of FE^p solutions gave rise to the development of reduced models, namely the Transformation Field Analysis (TFA, Dvorak and Benveniste, 1992) and its successor, the Nonuniform Transformation Field Analysis (NTFA, Michel and Suquet, 2003, 2004; Fritzen and Böhlke, 2010; Fritzen, 2011a). The key idea of the NTFA is the approximation of the inelastic strain field using a finite dimensional, spatially heterogeneous basis, denoted by inelastic modes. The number of degrees of freedom of the reduced problem refers to the number of mode activity coefficients which represent effective (macroscopic) internal variables. Usually this number is several magnitudes smaller than it is in the case of a full-field FEM simulation which reveals the computational efficiency of the method. In its original form, the NTFA (Michel and Suquet, 2003, 2004) determined the evolution of the mode activity based on a semi-phenomenological law. Besides giving good accuracy in comparison with full-field simulations, the simplicity of the evolution law is the bottleneck