



národní
úložiště
šedé
literatury

High performance computing in micromechanics

Blaheta, Radim
2014

Dostupný z <http://www.nusl.cz/ntk/nusl-170638>

Dílo je chráněno podle autorského zákona č. 121/2000 Sb.

Tento dokument byl stažen z Národního úložiště šedé literatury (NUŠL).

Datum stažení: 31.07.2024

Další dokumenty můžete najít prostřednictvím vyhledávacího rozhraní nusl.cz .

High performance computing in micromechanics

R. Blaheta, R. Hrtus, O. Jakl, J. Starý

Institute of Geonics AS CR, Ostrava

1 Introduction

By micromechanics we understand analysis of the macroscale response of materials through investigation of processes in their microstructure. Here by the macroscale, we mean the scale of applications, where we solve engineering problems involving materials like different metals and composites in aircraft design or rocks and concrete in a dam construction. Different applications are characterized by different characteristic size. At macroscale the materials mostly look as homogeneous or they are idealized as homogeneous or piecewise homogeneous. A substantial heterogeneity is hidden and appears only after more detailed zooming view into the material. This hidden heterogeneity can be called a microstructure. In metals it is created by crystals and grains, in composite materials by matrix and inclusions, in concrete by gravel and mortar or iron reinforcement etc. When the ratio between the characteristic dimensions on macro and microstructure subjects is sufficiently large, then we say that the scales are well separated. In this case, it is not possible to perform the macroscale analysis going into the microstructure details, but it is possible to analyse the macroscopic problems with the use of effective (homogenized) material properties, which are obtained by testing smaller samples of materials. In computational micromechanics, the testing of such samples means solution of boundary value problems on test domains involving the microstructure with loading provided by suitable boundary conditions.

We focus on X-ray CT image based micromechanics of geomaterials with the use of continuum mechanics and the finite element computation of the microscale strains and stresses, see [2]. This means that basic information about the microstructure is provided by analysing (segmentation) of 3D images of real samples. This information should be completed by information on local material properties, i.e. material properties of the individual material constituents.

There is a strong need for high performance parallel computing at several stages of the computational micromechanics, namely at

- analysis of CT scans,
- high resolution finite element solution of boundary value problems,
- solution of inverse problems for determination or calibration of local material properties.

In this contribution, we focus on the second point, i.e. solving the high resolution finite element systems with tens or hundreds degrees of freedom on available parallel computers at the Institute of Geonics and the IT4Innovations supercomputer centre in Ostrava. Following [3], we describe efficiency of the in-house GEM solvers exploiting the Schwarz domain decomposition method with aggregation by performing computational experiments on the above parallel computers. The solution of these systems is also necessary for building efficient solution methods for inverse material identification problems, see [4] and a further work in progress.

2 High resolution FEM systems and GEM solvers

In analysis of geocomposites (see [2]), the domain Ω is a cube with a relatively complicated microstructure. The FEM mesh is constructed on the basis of CT scans. As benchmarks, we shall use FEM systems arising from CT scanning of a coal-resin geocomposite at CT-lab of the Institute of Geonics. The characteristics of two benchmarks can be seen in Table 1.

Benchmark	Discretization	Size in DOF	Data size
GEOC-2s	$257 \times 257 \times 257$	50 923 779	8.5 GB
GEOC-2l	$257 \times 257 \times 1025$	203 100 675	33.5 GB

Table 1: Benchmarks representing microstructures of two geocomposite samples. Notation, applied discretization meshes and sizes of resulting linear systems.

The elastic response of a representative volume Ω is characterized by homogenized elasticity C or compliance S tensors ($S = C^{-1}$). The elasticity and compliance tensors are determined from the relations

$$C\langle\varepsilon\rangle = C\varepsilon_0 = \langle\sigma\rangle \text{ and } S\langle\sigma\rangle = S\sigma_0 = \langle\varepsilon\rangle, \quad (1)$$

respectively. Here $\langle\sigma\rangle$ and $\langle\varepsilon\rangle$ are volume averaged stresses and strains computed from the solution of elasticity problem

$$-\operatorname{div}(\sigma) = 0, \quad \sigma = C_m\varepsilon, \quad \varepsilon = (\nabla u + (\nabla u)^T)/2 \quad \text{in } \Omega, \quad (2)$$

with boundary conditions

$$u(x) = \varepsilon_0 \cdot x \text{ on } \partial\Omega \text{ and } \sigma \cdot n = \sigma_0 \cdot n \text{ on } \partial\Omega, \quad (3)$$

respectively. Above, σ and ε denote stress and strain in the microstructure, C_m is the variable local elasticity tensor, u and n denote the displacement and the unit normal, respectively. The use of pure Dirichlet and pure Neumann boundary conditions allows us to get a upper and lower bounds for the upscaled elasticity tensor, see e.g. [2].

By using the GEM software [1], the domain is discretized by linear tetrahedral finite elements. The arising systems are then solved by PCG method with a stabilization in the singular case (see [3]). The implementation in the GEM software uses two solvers:

GEM-DD is a solver implemented in the GEM software. It uses one-level additive Schwarz domain decomposition preconditioner with subproblems replaced by displacement decomposition incomplete factorization, see ref. in [3]. The resulting preconditioner is symmetric positive definite even for the singular case.

GEM-DD-CG solver differs in preconditioning, which is now a two-level Schwarz domain decomposition arising from the previous GEM-DD by additive involvement of a coarse problem correction. The coarse problem is created by a regular aggressive aggregation with 3 DOF's per aggregation. In singular case, the coarse problem is also singular with a smaller null space containing only the rigid shifts. The coarse problem is solved only approximately by inner (not stabilized) CG method with a lower solution accuracy - relative residual accuracy $\varepsilon_0 \leq 0.01$.

Note that in the computational experiments described in the next Section, we solve the problems with pure Neumann boundary conditions.

3 Parallel computers and computational experiments

The computational experiments are performed on two computers:

Enna - 64-core NUMA multiprocessor at the Institute of Geonics:

- eight octa-core Intel Xeon E7-8837/2.66 GHz processors
- 256 GB of DDR2 RAM
- CentOS 6.3, Intel Cluster Studio XE 2013, Trilinos 11.4.1

Anselm - multicomputer (cluster) with 207 compute nodes at the Supercomputing Center IT4Innovations. We employed the computing nodes equipped with:

- two octa-core Intel E5-2665/2.4 GHz processors
- 64 GB of memory and 500 GB of local disk capacity
- Infiniband QDR interconnection, fully non-blocking, fat-tree
- Bullx Linux OS (Red Hat family), Intel Parallel Studio XE 2013

Table 2 shows the timings of GEM solvers (without and with coarse grid problem applied) obtained for GEOC2s, i.e. a problem of more than 50 million DOF's, where the performance up to 64 processing elements on Enna and up to 128 processing elements on Anselm could be compared. The stopping criterion was $\|r\|/\|b\| \leq \varepsilon = 10^{-5}$ and the DD-CG solver made use of a coarse problem with aggregation factors $9 \times 9 \times 9$ (81 000 DOF's).

# Sd	Enna				Anselm			
	DD		DD-CG		DD		DD-CG	
	#It	T _{iter}	#It	T _{iter}	A/E	T _{iter}	A/E	T _{iter}
2	914	8461.2	437	3523.1	0.67	5644.2	0.79	2785.4
4	1129	4973.3	428	1923.6	0.59	3526.2	0.72	1383.4
8	1421	2942.5	416	922.9	0.82	2422.6	0.79	725.7
16	1655	1994.6	376	415.8	0.64	1325.8	0.84	348.7
32	1847	1923.5	329	348.3	0.42	798.3	0.56	194.8
64	2149	3074.9	295	505.9	0.20	620.8	0.23	117.6
128					n/a	515.7	n/a	107.1

Table 2: Timings of the GEOC2s benchmark achieved by the GEM solvers on the multiprocessor Enna and cluster Anselm: Iteration counts (#It), wall-clock time (in seconds) of the solution (T_{iter}) and the corresponding performance ratio Anselm/Enna (A/E) are provided for up to 128 subdomains (# Sd).

For greater number of subdomains, the results confirm the advantage of systems with distributed memory, when the multiprocessors in general suffer from the memory-processor bandwidth contention. Thus, while on Enna the scalability fades out at about 32 cores, the turning point on Anselm is around 128 processing elements, when the small size of subdomains deteriorates the computation/communication ratio.

In absolute figures, we were able to solve the benchmark 3–4 times faster on Anselm than on Enna. The advantage of Anselm is to be derived partially from the fact that its newer Intel Sandy Bridge CPU architecture as such outperforms Enna's Westmere one, in our application

# Sd	Enna						Anselm	
	$DD-9 \times 9 \times 9$		$DD-9 \times 9 \times 18$		$DD-9 \times 9 \times 27$		$DD-9 \times 9 \times 27$	
	# It	T_{iter}	# It	T_{iter}	# It	T_{iter}	# It	T_{iter}
4	751	13719.0	858	15737.6	997	18518.4	997	12671.4
8	690	6237.7	800	6960.8	917	8062.9	917	5803.9
16	585	2717.4	674	4010.6	777	4815.6	777	2576.6
32	585	2483.6	622	2923.8	708	3452.5	708	1157.5
64					627	3637.0	627	558.8
128							652	358.5
256							631	299.6
512							649	333.5

Table 3: Timings of the GEOC2l benchmark achieved by the GEM-DD-CG solver on the multiprocessor Enna and cluster Anselm: Iteration counts (#It) and wall-clock time (in seconds) for the solution time (T_{iter}) are provided now for different sizes of CG problem involved in computations and for various numbers of subdomains (# Sd).

by 20-40%, what can be estimated from the test up to 8 processing elements (one socket) when the processors work in similar conditions.

Table 3 reports computations with the largest benchmark GEOC2l (about 200 million DOF) and demonstrates the impact of the coarse grid size on the time of the solution. We can observe that very aggressive aggregation leads to the best results. We could confirm this observation on Anselm, where the best time in the Table 3 (299.6s with 256 processing elements and aggregation $9 \times 9 \times 27$) was surpassed by an experiment with the coarser aggregation $15 \times 15 \times 31$. The overall best GEOC2l solution time of 249.8s was achieved after 910 iterations on # Sd=512 subdomains (32 compute nodes employed).

A bit surprising decrease of the number of iterations with increasing number of subdomains (processors) as reported in the above Tables, especially for DD-CG, can be explained by the fact that smaller subdomain problems are solved more accurately in our implementation.

Acknowledgement: This work was supported by the European Regional Development Fund in the IT4Innovations Centre of Excellence project (CZ.1.05/1.1.00/02.0070).

References

- [1] R. Blaheta, O. Jakl, R. Kohut, J. Starý: *GEM – A platform for advanced mathematical geosimulations*. In: R. Wyrzykowski et al. (eds.), PPAM 2009, Part I, LNCS 6067, 266–275, 2010. and Parallel Computing PARA 2012, LNCS, Springer, accepted.
- [2] R. Blaheta, R. Kohut, A. Kolcun, K. Souček, L. Staš: *Micromechanics of geocomposites: CT images and FEM simulations*. In: EUROCK 2013, pp. 399–404, Balkema 2013.
- [3] R. Blaheta, O. Jakl, J. Starý, E. Turan: *Iterative solution of singular systems with applications*. Proceedings of the conference SNA 13, Institute of Geonics AS CR, 2013
- [4] R. Hrtus, J. Haslinger, R. Blaheta: *Identification problems with a priori given material interfaces*. WOFEX 2013