



Parallel PCG algorithms for voxel FEM elasticity systems

S. Margenov and Y. Vutov

Institute for Parallel Processing, Bulgarian Academy of Sciences
Acad. G. Bonchev, bl. 25A, 1113 Sofia, Bulgaria
margenov@parallel.bas.bg and *yavor@parallel.bas.bg*

Abstract. The presented comparative analysis concerns two parallel iterative solvers for large-scale linear systems related to μ FEM simulation of human bones. The benchmark problems represent the strongly heterogeneous structure of real bone specimens. The voxel data are obtained by a high resolution computer tomography. Non-conforming Rannacher-Turek finite elements are used for discretization of the considered problem of linear elasticity. It is well known that the preconditioned conjugate gradient method is the best tool for efficient solution of large-scale symmetric systems with sparse, positive definite matrices. Here, the performance of two parallel preconditioners is studied. Both are based on displacement decomposition. The first one uses modified incomplete Cholesky factorization MIC(0) and the other – algebraic multigrid. The comparative analysis is mostly based on the computing times to run the codes. The number of iterations for both preconditioners are also discussed.

Keywords: FEM, PCG, DD, MIC(0), AMG, parallel algorithms.

1 Introduction

This study is motivated by the development and tuning of robust iterative solution methods, algorithms and software tools for μ FE (micro finite element) simulation of human bones. A voxel representation of the bone structure based on micro computer tomography (CT) images is used to formulate the problem. The computational domain is a strongly heterogeneous composition of solid and fluid phases, see Fig. 2. The considered isotropic linear elasticity model is a current brick in the development of a toolkit for μ FE simulation of the bone microstructure. The implementation of a poroelasticity model is the next step in this project.

Non-conforming Rannacher-Turek FEs are used for discretization of the problem. The obtained linear system is large, with a sparse, symmetric and positive definite matrix. This implies the use of iterative solvers based on the preconditioned conjugate gradient (PCG) method [7]. The elasticity stiffness matrix has a coupled block structure corresponding to a separable displacement ordering of

the unknowns. Here, the performance of the following two basic preconditioning codes, incorporated in a displacement decomposition framework, is studied. The first one is the modified incomplete factorization, MIC(0), and the second is the algebraic multigrid, AMG. The MIC(0) code is developed in IPP-BAS, Sofia, while the AMG one is the BoomerAMG module of the software system HyPre developed at LLNL, Livermore. The comparative analysis is focused on the number of iterations and the related computing times for real-life large-scale problems.

The presented results are based on some recent studies of the authors addressed to the case of scalar elliptic problems [18]. The earlier considered FEM elliptic solvers are used here in the construction of efficient preconditioners for the coupled elasticity system.

The paper is organized as follows. In Section 2 we describe the Finite Element Method (FEM) setting of the problem. Then, in Section 3 MIC(0) and BoomerAMG preconditioning algorithms are presented. Section 4 is devoted to numerical experiments. First, experiments with constant coefficients on three parallel computers are shown. The next set of experiments illustrates the behavior of preconditioners on heterogeneous (voxel) problems with strong coefficient jumps.

2 Non-conforming FEM formulation of the problem

We consider the weak formulation of the linear elasticity problem in the form [1]: find $\mathbf{u} \in [H_E^1(\Omega)]^3 = \{\mathbf{v} \in [H^1(\Omega)]^3 : \mathbf{v}_{\Gamma_D} = \mathbf{u}_S\}$ such that

$$\int_{\Omega} [2\mu\varepsilon(\mathbf{u}) : \varepsilon(\mathbf{v}) + \lambda \operatorname{div} \mathbf{u} \operatorname{div} \mathbf{v}] d\Omega = \int_{\Omega} \mathbf{f}^t \mathbf{v} d\Omega + \int_{\Gamma_N} \mathbf{g}^t \mathbf{v} d\Gamma, \quad (1)$$

$\forall \mathbf{v} \in [H_0^1(\Omega)]^3 = \{\mathbf{v} = [H^1(\Omega)]^3 : \mathbf{v}_{\Gamma_D} = 0\}$, with the positive constants λ and μ of Lamé, the symmetric strains

$$\varepsilon(\mathbf{u}) := 0.5(\nabla \mathbf{u} + (\nabla \mathbf{u})^t),$$

the volume forces \mathbf{f} , and the boundary tractions \mathbf{g} , $\Gamma_N \cup \Gamma_D = \partial\Omega$, $|\Gamma_D| \neq \emptyset$. The Lamé coefficients are given by $\lambda = \frac{\nu E}{(1+\nu)(1-2\nu)}$, $\mu = \frac{E}{2(1+\nu)}$, where E stands for the modulus of elasticity, and $\nu \in (0, \frac{1}{2})$ is the Poisson ratio.

To obtain a stable saddle-point system one usually uses a mixed formulation for \mathbf{u} and $\operatorname{div} \mathbf{u}$. By the choice of piece-wise constant finite elements for the dual variable, it can be eliminated at the macroelement level, and thereafter a symmetric positive definite FEM system in primal unknowns (displacement) is obtained. This approach is known as *reduced and selective integration* (RSI) technique, see [3]. For the discretization of (1) we use nonconforming rotated trilinear elements of Rannacher-Turek [8].

Let $\Omega^H = w_1^H \times w_2^H \times w_3^H$ be a regular coarser decomposition of the domain $\Omega \subset \mathbb{R}^3$ into hexahedrons, and let the finer decomposition $\Omega^h = w_1^h \times w_2^h \times w_3^h$

be obtained by a regular refinement of each macro element $E \in \Omega^H$ into eight similar hexahedrons. The cube $\hat{e} = [-1, 1]^3$ is used as a reference element in the parametric definition of the rotated trilinear elements. For each $e \in \Omega^h$, let $\psi_e : \hat{e} \rightarrow e$ be the trilinear 1–1 transformation. Then the nodal basis functions are defined by the relations $\{\phi_i\}_{i=1}^6 = \{\hat{\phi}_i \circ \psi_e^{-1}\}_{i=1}^6$, where $\hat{\phi}_i \in \text{span}\{1, \xi_j, \xi_j^2 - \xi_{j+1}^2, j = 1, 2, 3\}$. One possible set of interpolation conditions are: $\hat{\phi}_i(b_{\Gamma^j}) = \delta_{ij}$, where δ_{ij} is the Kronecker symbol and b_{Γ^j} are the midpoints of the walls $\{\Gamma^j\}_{j=1}^6$ of \hat{e} . Then,

$$\begin{aligned}\hat{\phi}_1(x, y, z) &= (1 - 3x + 2x^2 - y^2 - z^2) / 6, \\ \hat{\phi}_2(x, y, z) &= (1 + 3x + 2x^2 - y^2 - z^2) / 6, \\ \hat{\phi}_3(x, y, z) &= (1 - x^2 - 3y + 2y^2 - z^2) / 6, \\ \hat{\phi}_4(x, y, z) &= (1 - x^2 + 3y + 2y^2 - z^2) / 6, \\ \hat{\phi}_5(x, y, z) &= (1 - x^2 - y^2 - 3z + 2z^2) / 6, \\ \hat{\phi}_6(x, y, z) &= (1 - x^2 - y^2 + 3z + 2z^2) / 6.\end{aligned}$$

The RSI FEM discretization reads as follows: find $\mathbf{u}^h \in V_E^h$ such that

$$\sum_{e \in \Omega^h} \int_e [2\mu \varepsilon^*(\mathbf{u}^h) : \varepsilon^*(\mathbf{v}^h) + \lambda \operatorname{div} \mathbf{u}^h \operatorname{div} \mathbf{v}^h] de = \int_{\Omega} \mathbf{f}^t \mathbf{v}^h d\Omega + \int_{\Gamma_N} \mathbf{g}^t \mathbf{v}^h d\Gamma, \quad (2)$$

$\forall \mathbf{v}^h \in V_0^h$, where $\varepsilon^*(\mathbf{u}) := \nabla \mathbf{u} - 0.5 I_L^{Q^H} [\nabla \mathbf{u} - (\nabla \mathbf{u})^t]$, V_0^h is the FEM space, satisfying (in nodalwise sense) homogeneous boundary conditions on Γ_D , the operator $I_L^{Q^H}$ denotes the L^2 -orthogonal projection onto Q^H , the space of piecewise constant functions on the coarser decomposition Ω^H of Ω . Then the standard computational procedure leads to the coupled system of linear equations

$$\begin{bmatrix} K_{11} & K_{12} & K_{13} \\ K_{21} & K_{22} & K_{23} \\ K_{31} & K_{32} & K_{33} \end{bmatrix} \begin{bmatrix} \mathbf{u}_h^1 \\ \mathbf{u}_h^2 \\ \mathbf{u}_h^3 \end{bmatrix} = \begin{bmatrix} \mathbf{f}_h^1 \\ \mathbf{f}_h^2 \\ \mathbf{f}_h^3 \end{bmatrix}. \quad (3)$$

Here the stiffness matrix K is written in block form corresponding to a separate displacements components ordering of the vector of nodal unknowns. Since K is sparse, symmetric and positive definite, we use the PCG method to solve the system (3).

3 Preconditioning algorithms

The PCG is known to be the best algorithm for solution of large systems of linear equations with symmetric and positive definite sparse matrices [7]. Crucial for its performance is the preconditioning technique used. Here we focus on two preconditioners based on the isotropic variant of the displacement decomposition (DD)[4,5]. We write the DD auxiliary matrix in the form

$$C_{DD} = \begin{bmatrix} A & & \\ & A & \\ & & A \end{bmatrix} \quad (4)$$

where A is the stiffness matrix corresponding to the bilinear form

$$a(u^h, v^h) = \sum_{e \in \Omega^h} \int_e E \left(\sum_{i=1}^3 \frac{\partial u^h}{\partial x_i} \frac{\partial v^h}{\partial x_i} \right) de. \quad (5)$$

Such approach is motivated by the second Korn's inequality, which holds for the RSI FEM discretization (2) under consideration. This means that the estimate

$$\kappa(C_{DD}^{-1}K) = O((1 - 2\nu)^{-1})$$

holds uniformly with respect to the mesh size parameter in the FEM discretization [6]. The first of the studied preconditioners is obtained by MIC(0) factorization of the blocks in (4). As an alternative, inner PCG iterations with BoomerAMG for A are used to approximate the DD block-diagonal matrix (4).

3.1 Parallel MIC(0) preconditioning

First, we give a brief introduction to the modified incomplete factorization [9], see also [11]. Our presentation at this point follows those from [4]. Let us rewrite the real $N \times N$ matrix $A = (a_{ij})$ in the form

$$A = D - L - L^T$$

where D is the diagonal and $(-L)$ is the strictly lower triangular part of A . Then we consider the approximate factorization of A which has the form:

$$C_{MIC(0)} = (X - L)X^{-1}(X - L)^T$$

with $X = \text{diag}(x_1, \dots, x_N)$ being the diagonal matrix determined by the condition of equal rowsums. We are interested in the case when $X > 0$ and thus $C_{MIC(0)}$ is positive definite for the purpose of preconditioning. If this holds, we speak about *stable* MIC(0) factorization. Concerning the stability of MIC(0), the following theorem holds [10,4].

Theorem 1. *Let $A = (a_{ij})$ be a symmetric real $N \times N$ matrix and let $A = D - L - L^T$ be the splitting of A . Let us assume that*

$$\begin{aligned} L &\geq 0, \\ Ae &\geq 0, \\ Ae + L^T e &> 0, \quad \mathbf{e} = (1, \dots, 1)^T \in R^N, \end{aligned}$$

i.e. that A is a weakly diagonally dominant with nonpositive offdiagonal entries and that $A + L^T = D - L$ is strictly diagonally dominant. Then the relation

$$x_i = a_{ii} - \sum_{k=1}^{i-1} \frac{a_{ik}}{x_k} \sum_{j=k+1}^N a_{kj} > 0$$

holds and the diagonal matrix $X = \text{diag}(x_1, \dots, x_N)$ defines stable MIC(0) factorization of A .

Remark 1. The presented numerical tests are performed using the perturbed version of MIC(0) algorithm, where the incomplete factorization is applied to the matrix $\tilde{A} = A + \tilde{D}$. The diagonal perturbation $\tilde{D} = \tilde{D}(\xi) = \text{diag}(\tilde{d}_1, \dots, \tilde{d}_N)$ is defined as follows:

$$\tilde{d}_i = \begin{cases} \xi a_{ii} & \text{if } a_{ii} \geq 2w_i \\ \xi^{1/2} a_{ii} & \text{if } a_{ii} < 2w_i \end{cases}$$

where $0 < \xi < 1$ is a properly chosen parameter, and $w_i = \sum_{j>i} -a_{ij}$.

The idea of our parallel algorithm is to apply the MIC(0) factorization on an auxiliary matrix B , which approximates A . The matrix B has a special block structure, which allows a scalable parallel implementation.

Following the standard FEM assembling procedure we write A in the form $A = \sum_{e \in \omega_h} L_e^T A_e L_e$, where A_e is the element stiffness matrix, L_e stands for the restriction mapping of the global vector of unknowns to the local one corresponding to the current element e . Let us consider the following approximation B_e of A_e :

$$A_e = \begin{bmatrix} a_{11} & a_{12} & a_{13} & a_{14} & a_{15} & a_{16} \\ a_{21} & a_{22} & a_{23} & a_{24} & a_{25} & a_{26} \\ a_{31} & a_{32} & a_{33} & a_{34} & a_{35} & a_{36} \\ a_{41} & a_{42} & a_{43} & a_{44} & a_{45} & a_{46} \\ a_{51} & a_{52} & a_{53} & a_{54} & a_{55} & a_{56} \\ a_{61} & a_{62} & a_{63} & a_{64} & a_{65} & a_{66} \end{bmatrix}, \quad B_e = \begin{bmatrix} b_{11} & a_{12} & a_{13} & a_{14} & a_{15} & a_{16} \\ a_{21} & b_{22} & a_{23} & a_{24} & a_{25} & a_{26} \\ a_{31} & a_{32} & b_{33} & 0 & 0 & 0 \\ a_{41} & a_{42} & 0 & b_{44} & 0 & 0 \\ a_{51} & a_{52} & 0 & 0 & b_{55} & 0 \\ a_{61} & a_{62} & 0 & 0 & 0 & b_{66} \end{bmatrix}.$$

The local numbering follows the pairs of the opposite nodes of the reference element. The diagonal entries of B_e are modified to hold the rowsum criteria. Assembling the locally defined matrices B_e we get the global matrix $B = \sum_{e \in \omega_h} L_e^T B_e L_e$. The condition number estimate $\kappa(B^{-1}A) \leq 3$ holds uniformly with respect to mesh parameter and possible coefficient jumps (see for the related analysis in [15,16]). This particular choice of B_e gives an important property of the matrix B – its diagonal blocks (corresponding to (x, y) cross sections) are diagonal matrices. This allows a scalable parallel implementation, which for the case of scalar elliptic problems is studied in [17,18]. The sparsity structure of the matrices A and B is illustrated by Fig. 1. Lexicographic node numbering is used.

This paper concerns with scalability of the parallel DD MIC(0) preconditioner

$$C_{DDMIC(0)} = \begin{bmatrix} C_{MIC(0)}(B) & & \\ & C_{MIC(0)}(B) & \\ & & C_{MIC(0)}(B) \end{bmatrix}.$$

3.2 BoomerAMG

BoomerAMG contains sequential and parallel implementations of algebraic multi-grid methods [12]. It can be used as a solver or as a preconditioner. Various different parallel coarsening techniques and relaxation schemes are available.

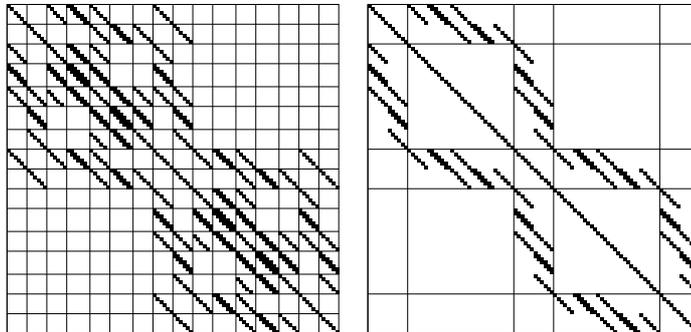


Fig. 1. Sparsity structure of the matrix A on the left and matrix B on the right, for the division of Ω into $2 \times 2 \times 6$ hexahedrons. Non-zero elements are drawn with small squares.

See [13,14] for a detailed description of the coarsening algorithms, the interpolation and numerical results.

Version 2.0.0 of the Hypr library was used for the performed tests.

The following coarsening techniques are available:

- the Cleary-Luby-Jones-Plassman (CLJP) coarsening,
- various variants of the classical Ruge-Stüben (RS) coarsening algorithm, and
- the Falgout coarsening which is a combination of CLJP and the classical RS coarsening algorithm.

The following relaxation techniques are available:

- Jacobi relaxation,
- hybrid Gauss-Seidel / Jacobi relaxation scheme,
- symmetric hybrid Gauss-Seidel / Jacobi relaxation scheme, and
- Gauss-Seidel relaxation.

The Falgout coarsening was used in the presented tests. A $V(1,1)$ -cycle with hybrid Gauss-Seidel smoothing is performed. The related AMG strength threshold is 0.5. The default coarsening with one inner iteration was set at the beginning. This variant proved to be very expensive, especially with respect to the consumed memory. Then we switched on the option of aggressive coarsening on the first two levels. As a result the memory consumption noticeably dropped. To get a reasonable convergence of the outer PCG iteration, the number of inner iterations were to be increased. Table 1 justifies our decision to use the aggressive coarsening in the further numerical tests. The following data are collected in this table: mesh size parameter n ; number of processors p ; number of AC levels (levels with aggressive coarsening where 0 stands for no aggressive coarsening); number of outer iterations It_{out} ; related inner iterations It_{inn} ; computation time

Table 1. BoomerAMG

n	p	AC levels	It_{out}	It_{inn}	$T[s]$	$M[MB]$
32	1	0	12	1	15.6	108
32	1	2	9	4	10.4	66
64	1	0	13	1	185.3	899
64	1	2	9	4	107.2	432
64	4	0	24	1	180.1	1225
64	4	2	10	4	58.5	664

T in seconds; required total memory M in megabytes. A more detailed description of the model test problem is given at the beginning of the next section. As one can see, despite the increased number of inner iterations, the computational times are strongly decreased and the parallel efficiency is improved if aggressive coarsening is used.

4 Comparative numerical tests

4.1 Scalability tests

Numerical tests with the considered two parallel algorithms and codes are present and analyzed in this section. The tests are run on three parallel platforms, referred to further as C1, C2 and C3. Platform C1 is an “IBM SP Cluster 1600” consisting of 64 p5-575 nodes interconnected with a pair of connections to the Federation HPS (High Performance Switch). Each p5-575 node contains 8 Power5 SMP processors at 1.9GHz and 16GB of RAM. The network bandwidth is 16Gb/s. Platform C2 is an IBM Linux Cluster 1350, made of 512 dual-core IBM X335 nodes. Each node contains 2 Xeon Pentium IV processors and 2GB of RAM. Nodes are interconnected with a 1Gb Myrinet network. Platform C3 is a “Cray XD1” cabinet, fully equipped with 72 2-way nodes, totaling in 144 AMD Opteron processors at 2.4GHz. Each node has 4GB of memory. The CPUs are interconnected with the Cray RaidArray network with a bandwidth of 5.6Gb/s.

The computational domain is the cube $[0, 1]^3$, where homogeneous Dirichlet boundary conditions are assumed at the bottom. The force $\|g\| = 1$ is acting on the top. The mesh is uniform. Here n stands for the number of subintervals in the fine grid of the RSI FEM discretization in each direction. The mechanical characteristics of the model problem are $E = 1$ and $\nu = 0.3$. The size of the resulting nonconforming FEM system is $N = 9n^2(n + 1)$. The number of processors p is increased proportionally with the problem size N . The stopping criterion in all considered tests is $(C^{-1}r^{N_{it}}, r^{N_{it}})/(C^{-1}r^0, r^0) < 10^{-6}$, where r^i is the current residual and C stands for the used preconditioner. Table 2 presents the time T in seconds, the number of iterations It (the outer ones for the AMG code), varying the preconditioners, the problem sizes and the platforms.

In a good agreement with the theory, the number of iterations for MIC(0) increases as $O(\sqrt{n})$, while the AMG iterations stay about the same. For the

Table 2. Parallel Tests I

n	N	p	C1				C2				C3			
			MIC(0)		AMG		MIC(0)		AMG		MIC(0)		AMG	
			$T[s]$	It										
64	2 396 160	1	136.6	115	150.1	9	83.7	115	84.0	9	83.9	115	115.1	9
128	19 021 824	8	202.0	163	195.6	10	172.1	163	229.8	10	127.8	163	152.6	10
256	151 584 768	64	355.6	230	261.4	10	464.1	230	430.0	10	328.2	230	307.1	10

smallest problem (N=2 396 160) MIC(0) clearly outperforms the AMG code. For the medium size (N= 19 021 824) the times are rather similar. However, for the largest problem (N=151 584 768) the advantage of AMG is well expressed.

4.2 Voxel analysis tests

The bone microstructure is a typical example of strongly heterogeneous media. In the presented tests, the computational domain is a composition of solid and fluid phases. The CT image is extracted from the dataset [2]. The voxel size is $37\mu\text{m}$. Each voxel corresponds to a macroelement from the RSI FEM discretization. The bone specimen is placed between two plates (see Fig. 2). The thickness of the plates is 1 voxel. The position of the bottom plate is fixed (homogeneous Dirichlet boundary conditions), and a force of $\|g\| = 1$ is uniformly distributed on the top one. This setting simulates a vertically loaded bone specimen. The density of the solid phase is an important characteristics of the bone. The density of the considered three specimens is respectively 25%, 19%, 18% for the cases $32 \times 32 \times 32$, $64 \times 64 \times 64$, $128 \times 128 \times 128$.

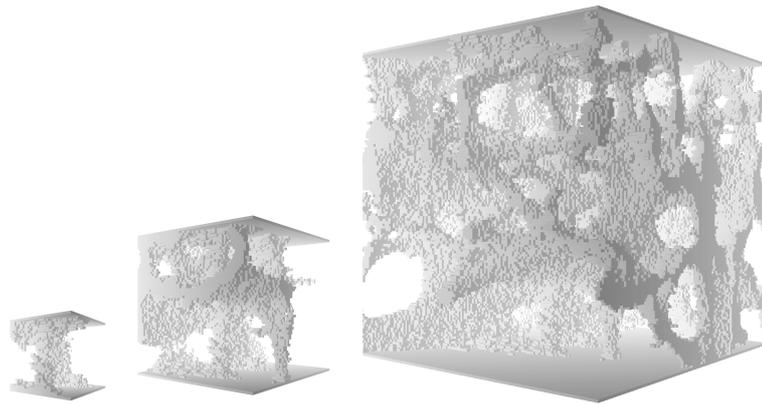


Fig. 2. Structure of the solid phase: $32 \times 32 \times 32$ - left, $64 \times 64 \times 64$ - middle, and $128 \times 128 \times 128$ - right.

The considered test problems are given by the following parameters: $E_p = 10$, $E_s = 1$, $E_f = \zeta \in \{0.1, 0.01, 0.001\}$, $\nu = 0.3$. Here, E_p is the elasticity modulus of the two plates, E_s stands for a scaled elasticity modulus of the solid phase, while E_f introduces varying coefficient jumps between solid and fluid phases.

The results presented in Table 3 are obtained on the platform C2. For the case of the biggest coefficient jumps ($\zeta=0.001$) and the biggest problem ($N=151\ 584\ 768$), outer PCG iteration with AMG preconditioner fails to converge within the specified time limit of 7200 seconds. This test was repeated with an increased number of inner iterations. The corresponding values in the table are obtained with $It_{in} = 6$.

Table 3. Parallel Tests II

		$\zeta = 0.1$				$\zeta = 0.01$				$\zeta = 0.001$			
n	p	MIC(0)		AMG		MIC(0)		AMG		MIC(0)		AMG	
		$T[s]$	It	$T[s]$	It	$T[s]$	It	$T[s]$	It	$T[s]$	It	$T[s]$	It
64	1	239.3	330	374.9	27	348.3	505	757.9	57	588.6	823	1040.5	78
128	8	833.2	708	681.0	25	975.5	830	1501.3	60	2166.7	1850	2908.9	107
256	64	2393.8	1237	945.4	25	3495.7	1831	2114.4	57	6025.8	3150	5520.1	114

What we see is that the number of iterations for the MIC(0) code increase more drastically with the problem size, than the case without jumps. AMG preconditioner manages to sustain the number of iterations for different problem sizes and fixed ζ , except for the case $\zeta = 0.001$, where slight increase is observed. For the smallest problem ($N=2\ 396\ 160$) MIC(0) outperforms the AMG code again for all variations of ζ . For the medium sized problem ($N=19\ 021\ 824$) MIC(0) code is faster for the cases of strong coefficient jumps ($\zeta = 0.01$ and $\zeta = 0.001$). For the largest problem ($N=151\ 584\ 768$) AMG is faster, but its advantage decreases with the rise of the coefficient jump. The reason for this behavior is that the AMG iterations are much more expensive than the MIC(0) ones. So, the increased number of iterations has much heavier impact on the computational time for the AMG code.

The general conclusion is that the studied codes provide a stable toolkit for computer simulation of the bone microstructure. Both approaches have their advantages depending on the size of the FEM systems and the level of heterogeneity of the bone specimens. The achieved parallel scalability well corresponds to the connectivity of the considered problems.

Acknowledgments. The partial support through EC INCO Grant BIS-21++ 016639/2005 and Bulgarian NSF Grant VU-MI-202/2006 is gratefully acknowledged. The parallel numerical tests are supported via EC Project HPC-EUROPA RII3-CT-2003-506079.

References

1. S. Margenov, *Displacement decomposition-MIC(0) preconditioning of linear elasticity non-conforming FEM problems*, 16th IMACS World Congress 2000, Lausanne, Proceedings, (2000), 107-4.
2. Gisela Beller, Markus Burkhart, Dieter Felsenberg, Wolfgang Gowin, Hans-Christian Hege, Bruno Koller, Steffen Prohaska, Peter I. Sapiro and Jesper S. Thomsen: Vertebral Body Data Set ESA29-99-L3, <http://bone3d.zib.de/data/2005/ESA29-99-L3/>
3. D. Malkus, T. Hughes. Mixed finite element methods. Reduced and selective integration techniques: an uniform concepts. *Comp. Meth. Appl. Mech. Eng.*,15: 63-81, 1978.
4. R. Blaheta: *Displacement Decomposition—incomplete factorization preconditioning techniques for linear elasticity problems*, Numer. Lin. Alg. Appl., **1** (1994), 107–126
5. O. Axelsson, I. Gustafsson. Iterative methods for the Navier equation of elasticity. *Comp. Meth. Appl. Mech. Engin.*, 15: 241-258, 1978.
6. Axelsson, O., On iterative solvers in structural mechanics; separate displacement orderings and mixed variable methods, *Math. Comput. Simulation* 50 (1999), no. 1–4, 11–30.
7. O. Axelsson: *Iterative Solution Methods*. Cambridge University Press, Cambridge, 1994.
8. R. Rannacher, S. Turek: *Simple nonconforming quadrilateral Stokes Element*, Numer. Methods Partial Differential Equations 8(2), 1992, 97–112.
9. I. Gustafsson: *Modified incomplete Cholesky (MIC) factorization*, In: D. J. Evans (ed.): *Preconditioning Methods; Theory and Applications*, Gordon and Breach, 1983, 265–293.
10. I. Gustafsson: An incomplete factorization preconditioning method based on modification of element matrices. *BIT* 35: 1, 86–100, 1996.
11. I. Gustafsson: *Stability and rate of convergence of modified incomplete Cholesky factorization methods*, Chalmers University of Technology, Report Nr. 79.02R, 1979.
12. J. W. Ruge and K. Stüben. *Algebraic multigrid (AMG)*. In S. F. McCormick, editor, *Multigrid Methods*, volume 3 of *Frontiers in Applied Mathematics*, SIAM, Philadelphia, PA, 1987, 73–130.
13. V. E. Henson and U. M. Yang. *BoomerAMG: a parallel algebraic multigrid solver and preconditioner*. *Applied Numerical Mathematics*, 41(5), 2002, 155–177. Also available as LLNL technical report UCRL-JC-141495.
14. U. M. Yang. *Parallel algebraic multigrid methods—high performance preconditioners*. In A. M. Bruaset and A. Tveito, editors, *Numerical Solution of Partial Differential Equations on Parallel Computers*, Springer-Verlag, 2005, 209–236. Also available as LLNL technical report UCRL-BOOK-208032.
15. P. Arbenz, S. Margenov: *Parallel MIC(0) preconditioning of 3D nonconforming FEM systems*, *Iterative Methods, Preconditioning and Numerical PDEs*, Proceedings of IMET Conference, Prague, 2004, 12–15.
16. P. Arbenz, S. Margenov, Y. Vutov: *Parallel MIC(0) Preconditioning of 3D Elliptic Problems Discretized by Rannacher-Turek Finite Elements*, *Comput. Math. Appl.* (to appear)
17. Y. Vutov: *Parallel Incomplete Factorization of 3D NC FEM Elliptic Systems*, *Numerical Methods and Applications*, Springer LNCS 4210, 2007, 114–121.
18. S. Margenov, Y. Vutov: *Preconditioning of Voxel FEM Elliptic Systems*, *TASK*, Quarterly 11, No 1–2, 2007, 117–128.