

Performance Analysis of Block AMG Preconditioning of Poroelasticity Equations

Nikola Kosturski¹, Svetozar Margenov¹, Peter Popov^{1,2}, Nikola Simeonov²,
and Yavor Vutov¹

¹ Institute of Information and Communication Technologies,
Bulgarian Academy of Sciences, Sofia, Bulgaria

² PPResearch Ltd, Sofia, Bulgaria

Abstract. The goal of this study is to develop, analyze, and implement efficient numerical algorithms for equations of linear poroelasticity, a macroscopically diphasic description of coupled flow and mechanics. We suppose that the solid phase is governed by the linearized constitutive relationship of Hooke's law. Assuming in addition a quasi-steady regime of the fluid structure interaction, the media is described by the Biot's system of equations for the unknown displacements and pressure (\mathbf{u}, p) . A mixed Finite Element Method (FEM) is applied for discretization. Linear conforming elements are used for the displacements. Following the approach of Arnold-Brezzi, non-conforming FEM approximation is applied for the pressure where bubble terms are added to guarantee a local mass conservation. Block-diagonal preconditioners are used for iterative solution of the arising saddle-point linear algebraic system. The BiCGStab and GMRES are the basic iterative schemes, while algebraic multigrid (AMG) is utilized for approximation of the diagonal blocks. The HYPRE implementations of BiCGStab, GMRES and AMG (BoomerAMG, [6]) are used in the presented numerical tests. The aim of the performance analysis is to improve both: (i) the convergence rate of the solvers measured by the iteration counts, and (ii) the CPU time to solve the problem. The reported results demonstrate some advantages of GMRES for the considered real-life, large-scale, and strongly heterogeneous test problems. Significant improvement is observed due to tuning of the BoomerAMG settings.

1 Introduction

In classical linear poroelasticity it is assumed that the solid is governed by the constitutive relationship $\sigma = \mathcal{L}\mathbf{e}$, where σ is the stress tensor, $\mathbf{e}(\mathbf{u}) = \frac{1}{2}(\nabla\mathbf{u} + \nabla\mathbf{u}^T)$ is the strain tensor and \mathcal{L} stands for the elasticity tensor. Then the media is described by the Biot law (c.f. [4, 9], see also [10]):

$$\nabla \cdot (\mathcal{L}\mathbf{e}(\mathbf{u}) - \mathbf{A}p_0) = \mathbf{0}, \quad (1)$$

$$\nabla \cdot (\mathbf{K}\nabla p) = \nabla \cdot \frac{\partial \mathbf{u}}{\partial t} + \mathbf{A} : \mathbf{e} \left(\frac{\partial \mathbf{u}}{\partial t} \right) + \beta \frac{\partial p}{\partial t}. \quad (2)$$

Here p and \mathbf{u} are the unknown pressure and displacement vector. \mathbf{K} stands for the permeability tensor. The Biot coefficient (tensor) \mathbf{A} takes into account the contribution of the fluid pressure p into the momentum equation (1), as well as the pore volume change term $\mathbf{A} : \mathbf{e} \left(\frac{\partial \mathbf{u}}{\partial t} \right)$ in the balance of mass (1), due to the displacements \mathbf{u} . The pore volume change in (2) due to p is captured by β , the coefficient of apparent rock compressibility due macroscopic fluid pressure.

The mesh methods provide computational technology for efficient discretization of the problem (1-2). Among others, we would mention the Galerkin finite element method (FEM) and the mixed FEM. The choice of method depends on the features of the considered class of problems. Here, conforming linear FEs are applied for approximation of the displacements in the elasticity terms of (1-2). For applications related to flows in highly heterogeneous porous media, the mixed finite element methods have proven to be accurate and locally mass conservative. While applying the mixed FEM to fluid subproblem, the continuity of the velocity normal to the boundary between two adjacent finite elements could be enforced by Lagrange multipliers. The relationship between the mixed and non-conforming FEM has been studied and simplified for various finite element spaces (see, e.g. [2]). In [3] Arnold and Brezzi have demonstrated that after the elimination of the unknowns representing the pressure and the velocity from the algebraic system the resulting Schur system for the Lagrange multipliers is equivalent to a discretization by Galerkin method using linear non-conforming finite elements. Namely, in [3] it is shown that the lowest-order Raviart-Thomas mixed finite element approximations are equivalent to the usual Crouzeix-Raviart non-conforming linear finite element approximations when the non-conforming space is augmented with quadratic bubbles. We use such kind of augmented Crouzeix-Raviart linear elements for approximation of the pressure. Implicit backward Euler method is applied for time discretization.

The rest of the paper is devoted to the solution of the linear algebraic systems arising at each time step. It is organized as follows. In Section 2, we consider the preconditioning algorithms for the related saddle-point problems. Section 3 contains the key results. A performance analysis based on tuning of a set of preconditioning parameters is presented here. Some concluding remarks are given at the end.

2 Preconditioning

The saddle-point poroelasticity problem corresponding to a displacements-pressure two by two block splitting (see Fig. 1 (b)) is written in the form:

$$A^{Biot} = \begin{bmatrix} A_{\mathbf{u}\mathbf{u}} & A_{\mathbf{u}p} \\ A_{p\mathbf{u}} & A_{pp} \end{bmatrix}$$

To get a symmetric matrix A^{Biot} , the pressure equation (2) is multiplied by -1 .

For large-scale FEM systems, the advantages of the iterative solution methods are well known. There are efficient Krylov subspace methods designed for




Fig. 1. (a) Example of strongly anisotropic unstructured FEM mesh in reservoir simulation; (b) Structure of the nonzero entries of the saddle-point matrix A^{Biot}

solving symmetric indefinite problems. In this study, the behaviour of the bi-conjugate gradient stabilized method BiCGStab and the generalized minimal residual method GMRES are studied. The potentially increasing memory consumption of GMRES is controlled with restarts.

A crucial ingredient for success of the Krylov subspace methods is the preconditioning. In this work, we use a block-diagonal preconditioning. Some related results for poroelasticity FEM systems are available in some recent papers (c.f. e.g. [1, 5] and the references there in), where the media is either homogeneous or not strongly heterogeneous. The robustness for heterogeneous problems of high-contrast is a hot topic. The derivation of uniform inf-sup estimate with respect to the coefficient jumps (c.f. e.g. [7]) is still a challenging issue for poroelasticity FEM systems.

The BoomerAMG preconditioner [6] is used as approximation of the diagonal blocks $A_{\mathbf{uu}}$ and A_{pp} . The resulting linear preconditioner of the Biot matrix is in the form:

$$C^{Biot} = \begin{bmatrix} A_{\mathbf{uu}}^{AMG} & \\ & A_{pp}^{AMG} \end{bmatrix}$$

The following 3D Test Problems (TP) concerning simulation of flows in deformable porous media as appear in petroleum engineering (see Fig. 1(a)) are considered in the next section: (i) TP1: $N = N_{\mathbf{u}} + N_p = 3 \times 2\,291 + 9\,592 = 16\,465$; the media is modestly heterogeneous; (ii) TP2: $N = N_{\mathbf{u}} + N_p = 3 \times 150\,871 + 186\,669 = 639\,282$; the media is strongly heterogeneous; the mesh is strongly anisotropic. Here N stands for the total number of degrees of freedom, while $N_{\mathbf{u}}$ and N_p stand for the related numbers of displacements and pressures. Unstructured grids are used in both, TP1 and TP2.

The large condition number of TP2 is due to the following complementary factors: the elasticity modulus range is $E \in (0.5, 30)$ [GPa]; the Poisson ratio range is $\nu \in (0.3, 0.48)$; the permeability variation is of order of 5×10^5 ; the mesh anisotropy ratio of order of 6×10^3 .

The distribution of nonzero entries of the matrix corresponding to TP2 is presented in Fig. 1(b). The structure of A^{Biot} corresponds to the following num-

Table 1. Performance of BiCGStab iterative solver

ϵ	TP1		TP2	
	N_{it}^{BiCG}	T[s]	N_{it}^{BiCG}	T[s]
10^{-6}	7	0.23	18	26.48
10^{-9}	10	0.28	63	79.67
10^{-12}	14	0.36	430	517.30

bering: (i) pointwise ordering in A_{uu} (we don't assume separable displacements);
(ii) the unknowns corresponding to the bubbles in A_{pp} are at the end.

3 Performance Analysis: Tuning of the BoomerAMG Parameters

The numerical tests are performed on a 3.4 GHz Intel Core i7 CPU. The following notations are used: ϵ - relative stopping criteria for both BiCGStab and GMRES; N_{it}^{BiCG} - number of BiCGStab iterations; N_{it}^{GMRES} - number of GMRES iterations; T[s] - CPU time in seconds. The performance analysis is started with the default settings of BoomerAMG: Falgout coarsening, hybrid symmetric Gauss-Seidel relaxation, classical modified interpolation, and a Strong Threshold equal to 0.75. The results for BiCGStab are given in Table 1.

The influence of the size of Krylov subspace before restart, $Kdim$, is examined additionally when the GMRES performance is studied. In Table 2, we see how the number of iterations N_{it}^{GMRES} decreases with the increase of $Kdim$. The CPU time is almost always smaller for largest values of $Kdim$. GMRES is faster, even though the count of BiCG iterations is smaller.

In the following performance analysis, the numerical tests are only for the larger problem TP2, where the heterogeneity and mesh anisotropy are very strong. Comprehensive tests for BiCGStab ($\epsilon = 10^{-6}$) are performed at the next

Table 2. Performance of GMRES iterative solver

TP	ϵ	Kdim	4	8	16	32	64	128	256	512
1	10^{-6}	N_{it}^{GMRES}	24	13	11					
		T[s]	0.38	0.25	0.22					
1	10^{-9}	N_{it}^{GMRES}	39	21	16					
		T[s]	0.53	0.32	0.28					
1	10^{-12}	N_{it}^{GMRES}	47	28	24	21				
		T[s]	0.62	0.40	0.35	0.31				
2	10^{-6}	N_{it}^{GMRES}		62	44	28				
		T[s]		46.75	33.28	22.81				
2	10^{-9}	N_{it}^{GMRES}		151	128	92	88	84		
		T[s]		105.94	86.84	63.37	61.69	60.44		
2	10^{-12}	N_{it}^{GMRES}			893	617	468	372	362	351
		T[s]			574.08	396.13	307.16	259.86	275.30	302.09

Table 3. Tuning of the Coarsening: TP2, $\epsilon = 10^{-6}$, $Kdim = 32$; Relaxation - 1; Interpolation - 0

		BiCGStab		GMRES	
Coarsening	N_{it}	$T[s]$	N_{it}	$T[s]$	
21	17	17.02	28	15.80	
22	17	17.08	28	15.83	
3	17	17.16	28	15.87	
0	17	18.44	28	16.44	
6	17	19.05	29	16.73	

step, varying the following parameters: (i) Coarsening: 0 - CLJP-coarsening; 3 - Ruge-Stueben coarsening; 6 - Falgout coarsening; 21 - CGC coarsening; 22 - CGC-E coarsening; (ii) Relaxation: 1 - Gauss-Seidel, sequential; 3 - hybrid Gauss-Seidel forward solve; 4 - hybrid Gauss-Seidel backward solve; 5 - hybrid chaotic Gauss-Seidel; 6 - hybrid symmetric Gauss-Seidel; (iii) Interpolation: 0 - classical modied interpolation; 4 - multipass interpolation; 6 - extended+i interpolation; 7 - extended+i (if no common C neighbor) interpolation; 8 - standard interpolation; 12 - FF interpolation; 13 - FF1 interpolation; 14 - extended interpolation. The best result is obtained for the setting: Coarsening - 21; Relaxation - 1; Interpolation - 0. This variant is selected as default for the next parameter by parameter tunings. The comparative results for both, BiCGStab and GMRES ($Kdim = 32$), are shown in Tables 3-5.

Now, we analyze the influence of the so called Strong Threshold which (according to HYPRE documentation) is to be chosen in the interval $(0, 1)$, depending on the particular problem. The results are given in Table 6. What we observe is the monotone increasing of the number of iterations, pursued with a monotone decreasing of the time for both, BiCGStab and GMRES.

The last step of the presented performance analysis is devoted to tuning the BoomerAMG parameters separately for each of the blocks A_{uu} and A_{pp} where a stronger stopping criteria of $\epsilon = 10^{-12}$ is applied. Some selected best settings are given in the Tables 7-8.

The final result of this experimental study is given in Table 9.

Table 4. Tuning of the Relaxation: TP2, $\epsilon = 10^{-6}$, $Kdim = 32$; Coarsening - 21; Interpolation - 0

		BiCGStab		GMRES	
Relaxation	N_{it}	$T[s]$	N_{it}	$T[s]$	
5	17	17.05	28	15.87	
3	17	17.15	28	15.98	
1	17	17.21	28	15.85	
4	17	19.16	28	17.57	
6	18	25.98	27	21.86	

Table 5. Tuning of the Interpolation: TP2, $\epsilon = 10^{-6}$, $Kdim = 32$; Coarsening - 21; Relaxation - 1

Interpolation	BiCGStab		GMRES	
	N_{it}	T[s]	N_{it}	T[s]
7	17	16.98	28	15.90
4	17	17.05	28	15.95
12	17	17.08	28	15.94
0	17	17.09	28	15.77
13	17	17.16	28	15.87
14	17	17.29	28	16.01
6	20	19.52	28	15.98
8	21	20.35	28	16.14

Table 6. Tuning of the Strong Threshold: TP2, $\epsilon = 10^{-6}$, $Kdim = 32$; Coarsening - 21; Relaxation - 1

ST	BiCGStab		GMRES	
	N_{it}	T[s]	N_{it}	T[s]
0.05	16	50.98	21	44.14
0.10	15	42.79	20	36.52
0.15	14	37.47	20	32.50
0.20	13	33.00	21	30.43
0.30	14	29.35	21	25.86
0.40	16	27.30	23	23.09
0.50	16	23.45	23	19.63
0.60	16	19.99	25	17.98
0.70	18	19.00	28	16.85
0.75	17	17.00	28	15.91
0.80	18	16.81	29	15.31
0.85	21	17.83	30	14.77
0.90	20	16.35	31	14.55
0.95	20	15.59	31	13.82

Table 7. Tuning of the BoomerAMG parameters for the block A_{uu} : TP2, $\epsilon = 10^{-12}$

Coarsening	Relaxation	Interpolation	N_{it}	T[s]
21	5	12	50	21.37
21	5	0	50	21.42
3	5	12	50	21.42
22	5	4	50	21.42
22	1	13	50	21.44
3	5	7	50	21.44

Table 8. Tuning of the BoomerAMG parameters for the block A_{pp} : TP2, $\epsilon = 10^{-12}$

Coarsening	Relaxation	Interpolation	N_{it}	T[s]
0	6	13	789	33.37
6	6	6	794	33.39
6	6	12	794	33.39
6	6	0	794	33.39
6	6	14	794	33.40
0	6	12	789	33.43

Table 9. Behavior of the solvers of the coupled system after tuning the BoomerAMG parameters: TP2, Kdim=32

ϵ	BiCGStab		GMRES	
	N_{it}	T[s]	N_{it}	T[s]
10^{-6}	14	15.91	22	13.94
10^{-9}	54	45.27	68	34.56
10^{-12}	279	213.36	282	170.98

4 Concluding Remarks

Block-diagonal preconditioning of the mixed FEM Biot system is studied. One commonly used technique is based on inner iterations for the related elliptic blocks. Let us note that the related preconditioners are not linear. Numerical tests illustrating the efficiency of this approach for modestly heterogeneous problems can be found in [1, 5]. Our study is focussed on problems with strong heterogeneity and strong mesh anisotropy (see the related details for TP2). For this class of problems, we don't observe any advantages of inner iterations. This is the reason to concentrate on linear preconditioners where AMG approximation of the diagonal blocks is used. The presented results are of strongly expressed experimental nature. The performance analysis shows a serious potential for improvement of the computational efficiency. Due to tuning of the BoomerAMG parameters, the achieved decrease of the CPU times T[s] for TP2 for $\epsilon = 10^{-12}$ are as follows: (i) BiCGStab - from 517.30 to 213.36, that is a reduction factor of 2.42; (ii) GMRES (Kdim=32) - from 396.13 to 170.98, or a reduction factor of 2.32. The better performance of GMRES is clearly visible. In this respect we have to remember that the required memory increases with Kdim which could be a restriction for more large scale applications.

Acknowledgments

The research is partly supported by the project AComIn "Advanced Computing for Innovation", grant 316087, funded by the FP7 Capacity Program.

References

1. P. Arbenz, E. Turan, Preconditioning for Large Scale Micro Finite Element Analyses of 3D Poroelasticity, Springer LNCS 7782 (2013), 361–374.
2. T. Arbogast and Z. Chen, On the implementation of mixed methods as non-conforming methods for second order elliptic problems, *Math. Comp.* 64 (1995), 943–972.
3. D.N. Arnold and F. Brezzi, Mixed and nonconforming finite element methods: implementation, postprocessing and error estimates, *RAIRO, Model. Math. Anal. Numer.* 19 (1985), 7–32.
4. M.A. Biot, General theory of three dimensional consolidation, *J. Appl. Phys.* 12 (1941), 155–164.
5. J.B. Haga, H. Osnes, H.P. Langtangen, A parallel preconditioner for large-scale poroelasticity with highly heterogeneous material parameters, *Comput. Geosci.* 16 (2012), 723–734.
6. V. E. Henson, U. M. Yang, BoomerAMG: a parallel algebraic multigrid solver and preconditioner, *Appl. Num. Math.* 41(5) (2002), 155–177.
7. J. Kraus, R. Lazarov, M. Limbery, S. Margenov, L. Zikatanov, Preconditioning of weighted $H(\text{div})$ -norm and applications to numerical simulation of highly heterogeneous media, Cornell University Library, arXiv:1406.4455, 2014.
8. J. Kraus, S. Margenov, Robust Algebraic Multilevel Methods and Algorithms, Radon Series on Computational and Applied Mathematics 5, de Gruyter, 2009
9. C.K. Lee, C.C. Mei, Re-examination of the equations of poroelasticity, *Int. J. Eng. Sci.* 35 (1997), 329–352.
10. P. Popov, Upscaling of Deformable Porous Media with Applications to Bone Modelling, Annual Meeting of the Bulgarian Section of SIAM, BGSIAM09 Proceedings, 2010, 105–110.