Numerical Experiments with an Overlapping Additive Schwarz Solver for 3-D Parallel Reservoir Simulation

Luca F. Pavarino
Marcelo Ramé

CRPC-TR93373
June 1993

Center for Research on Parallel Computation
Rice University
P.O. Box 1892
Houston, TX 77251-1892

Numerical experiments with an overlapping additive Schwarz solver for 3-D parallel reservoir simulation

Luca F. Pavarino and Marcelo Ramé *

Abstract

We present some numerical experiments with an overlapping additive Schwarz method for the mixed finite element discretization of second order elliptic problems in three dimensions, arising from flow models in reservoir simulation. These problems are characterized by large uncorrelated variations in the coefficients of the elliptic operator, which make the problems very ill- conditioned. The results confirm the theoretical bound on the condition number of the iteration operator and show the advantage of domain decomposition preconditioning as opposed to a simpler but less robust diagonal preconditioner.

1 Introduction

Domain decomposition methods are a major area of contemporary research in numerical analysis of partial differential equations. The discretization of elliptic problems arising in reservoir simulation and, more generally, in simulation of flow in porous media, are characterized by large, sparse and very ill-conditioned linear systems. This is mainly caused by rapid and large changes in the flow coefficients of the ellipitic operator (in typical rock, changes in permeability of three to five orders of magnitude over a few feet are not uncommon) as well as by the large scale of the problem. In this paper, we solve these linear systems with an overlapping additive Schwarz method (see Dryja and Widlund [6]), which has been proven to be

*Department of Computational and Applied Mathematics, Rice University, Houston, TX 77251. email: pavarino@rice.edu, marcelo@rice.edu.

very parallelizable, scalable and robust for three dimensional Galerkin problems, see Gropp and Smith [11]. Here, we study the method applied to mixed finite element problems. We refer to Cowsar [3], Ewing and Wang [8] and Mathew [12], [13], [14] for the mathematical analysis of this approach.

This paper is organized as follows. In the next section, we introduce the elliptic problem and its mixed finite element discretization. In Section 3, we briefly recall the definitions and basic results for overlapping additive Schwarz methods and then we introduce the algorithm implemented in this work. Some implementation details are given in Section 4, while the numerical results are presented in Section 5.

2 The Elliptic Problem

Let Ω be a bounded domain in R^3 with piecewise smooth boundary Γ. Given a 3×3 uniformly positive definite matrix $A(x)$ (which represents the flow coefficient tensor) and $f \in L^2(\Omega)$, we consider the elliptic problem

$$\begin{aligned} -\nabla \cdot A(x)\nabla p &= f \quad \text{in} \quad \Omega \\ \nabla p \cdot n &= 0 \quad \text{on} \quad \Gamma. \end{aligned} \tag{1}$$

Introducing the new vector unknown $\mathbf{u} = -A(x)\nabla p$, we can write (1) as

$$\begin{aligned} A^{-1}\mathbf{u} &= -\nabla p \quad \text{in} \quad \Omega \\ \nabla \cdot \mathbf{u} &= f \quad \text{in} \quad \Omega \\ \nabla p \cdot n &= 0 \quad \text{on} \quad \Gamma. \end{aligned} \tag{2}$$

The variational formulation of this problem consists in finding $(\mathbf{u}, p) \in H_0(div;\Omega) \times L^2(\Omega)$ such that

$$\int_\Omega A^{-1}\mathbf{u} \cdot \mathbf{v} dx - \int_\Omega p \nabla \cdot \mathbf{v} dx = 0, \qquad \forall \mathbf{v} \in H_0(div;\Omega), \tag{3}$$

$$\int_\Omega \nabla \cdot \mathbf{u} q dx = \int_\Omega f q dx, \qquad \forall q \in L^2(\Omega), \tag{4}$$

where $H_0(div;\Omega)$ is the kernel of the normal trace mapping of $H(div;\Omega)$ into $L^2(\Gamma)$. In order to discretize (3) and (4), we introduce a triangulation $\mathcal{T}_h$ of Ω into elements of size h satisfying the usual regularity requirements. We also consider a coarse triangulation of Ω into nonoverlapping subdomains Ω_i of size H, each consisting of a union of elements of $\mathcal{T}_h$. A standard mixed finite element approximation of (3) and (4) is obtained by introducing finite dimensional subspaces $W_h(\Omega) \subset H_0(div;\Omega)$

and $V_h(\Omega) \subset L^2(\Omega)$ associated to the triangulation $\mathcal{T}_h$: find $(\mathbf{u}_h, p_h) \in W_h(\Omega) \times V_h(\Omega)$ such that

$$\int_\Omega A^{-1}\mathbf{u}_h \cdot \mathbf{v}_h dx - \int_\Omega p_h \nabla \cdot \mathbf{v}_h dx = 0, \qquad \forall \mathbf{v}_h \in W_h(\Omega), \tag{5}$$

$$\int_\Omega \nabla \cdot \mathbf{u}_h q_h dx = \int_\Omega f q_h dx, \qquad \forall q_h \in V_h(\Omega). \tag{6}$$

We assume that this problem is well posed, i.e. the Babuska-Brezzi inf-sup condition holds, see Brezzi [2]. For a review of the most important mixed finite element spaces, see Roberts and Thomas [15]. We eliminate the velocity unknown $\mathbf{u}_h$ by introducing a discrete gradient operator $\nabla_h : V_h(\Omega) \rightarrow W_h(\Omega)$ such that

$$\int_\Omega A^{-1}(\nabla_h[q_h]) \cdot \mathbf{v}_h dx = -\int_\Omega q_h \nabla \cdot \mathbf{v}_h dx, \qquad \forall \mathbf{v}_h \in W_h(\Omega). \tag{7}$$

Problem (5) and (6) is then equivalent to the problem for the pressure unknown p_h

$$a(p_h, q_h) = \int_\Omega f q_h dx, \qquad \forall q_h \in V_h(\Omega), \tag{8}$$

where

$$a(p_h, q_h) = \int_\Omega A^{-1}\nabla_h[p_h]\nabla_h[q_h]dx = -\int_\Omega q_h \nabla \cdot (\nabla_h[p_h])dx.$$

We recall that by computing the integrals in (7) by special quadrature rules, the resulting linear system is equivalent to a cell- centered finite difference discretization, see Weiser and Wheeler [17]. We solve this linear system by an iterative method using an overlapping domain decomposition preconditioner of additive Schwarz type.

3 Additive Schwarz Methods

3.1 Abstract theory

For a detailed presentation of additive, multiplicative and hybrid Schwarz methods, see Dryja and Widlund [7] and Dryja, Smith and Widlund [5] and the references therein. Let V be a finite dimensional Hilbert space and $a(\cdot,\cdot) : V \times V \rightarrow R$ a selfadjoint, elliptic and bounded bilinear form. We are solving the problem:
find $u \in V$ such that

$$a(u, v) = f(v), \qquad \forall v \in V. \tag{9}$$

The space V can be decomposed into a sum (not necessarily direct) of $N + 1$ subspaces

$$V = V_0 + V_1 + \cdots + V_N,$$

where the first space V_0 is related to a special coarse discretization of the problem. Let $b_i(\cdot,\cdot) : V_i \times V_i \to R$ be inner products and $T_i : V \to V_i$ be operators defined by

$$b_i(T_i u, v) = a(u, v), \qquad \forall v \in V_i.$$

If $b_i(u, v) = a(u, v)$, then $T_i = P_i$, the a-orthogonal projection onto V_i. This choice corresponds to the use of exact solvers in the algorithm. The additive Schwarz operator is defined as

$$T = T_0 + T_1 + \cdots + T_N.$$

The original problem (9) is replaced by the preconditioned problem

$$Tu = g, \tag{10}$$

where the right hand side $g = \sum_{i=0}^{N} g_i = \sum_{i=0}^{N} T_i u$ is constructed by solving

$$b_i(g, v) = a(u, v) = f(v), \qquad \forall v \in V_i.$$

This linear system is solved iteratively (with a method that does not require the explicit construction of the stiffness matrix, but only its application to a given vector) and the iteration is usually accelerated by a Krilov space method: if the original problem is symmetric and positive definite, we can use the conjugate gradient method, otherwise we can use GMRES or other methods for general systems, see Freund, Golub and Nachtigal [9]. Dryja and Widlund proved the following basic

Theorem 1 *Let*
(i) there exists a constant C_0 such that for all $u \in V$ there exists a decomposition $u = \sum_{i=0}^{N} u_i$, $u_i \in V_i$, such that

$$\sum_{i=0}^{N} b_i(u_i, u_i) \leq C_0^2 a(u, u);$$

(ii) there exists a constant ω such that for $i = 0, 1, \cdots, N$

$$a(u, u) \leq \omega b_i(u, u), \qquad \forall u \in V_i;$$

(iii) there exist constants $\epsilon_{ij}, i, j = 1, \cdots, N$, such that

$$a(u_i, u_j) \leq \epsilon_{ij} a(u_i, u_i)^{1/2} a(u_j, u_j)^{1/2}, \qquad \forall u_i \in V_i,\ \forall u_j \in V_j.$$

Then T is invertible and

$$C_0^{-2} a(u, u) \leq a(Tu, u) \leq (\rho(\epsilon) + 1)\omega a(u, u), \qquad \forall u \in V.$$

Here $\rho(\epsilon)$ is the spectral radius of the matrix $\mathcal{E} = \{\epsilon_{ij}\}_{ij=1}^{N}$.

In the applications, it is particularly important the case where (9) is the discretization of an elliptic problem over a domain Ω decomposed into subdomains Ω_i. In the classic overlapping additive Schwarz method, each Ω_i is extended a certain number of elements beyond the boundary to a larger subdomain Ω_i'. The number of elements used in this extension determines the overlap of the new decomposition $\{\Omega_i'\}$, which is measured by

$$\delta = \min_i dist(\partial\Omega_i \backslash \partial\Omega, \partial\Omega_i' \backslash \partial\Omega).$$

The space decomposition in the Galerkin case is given by

- V_0 = space of trilinear basis functions defined on the coarse mesh and satisfying the given boundary conditions;
- $V_i = V_h \bigcap H_0^1(\Omega_i'), i = 1, \cdots, N.$

Dryja and Widlund [6] proved the following

Theorem 2 *When exact solvers are used for the subproblems, the condition number of the additive Schwarz method satisfies*

$$\kappa(P) \leq C(1 + \frac{H}{\delta}),$$

where C is independent of H, h and δ.

Cowsar [3] proved that this result is valid also for problems discretized with mixed and hybrid finite element methods.

3.2 The algorithm

We now describe an additive Schwarz method with minimal overlap for the discrete problem (8). For each subdomain Ω_i, let Ω_i' be the extension with minimal overlap ($\delta = h$). The coarse and local subspaces are

- $V_0 = \{\phi \in V_h | \phi =$ interpolant at the V_h nodes of the trilinear basis functions of $V_H\}$;
- $V_i = \{\phi \in V_h | supp(\phi) \subset \Omega_i'\}$.

Applying the results presented in the previous section, in particular Theorem 7.2 in Cowsar [3], we have

Theorem 3 *When exact solvers are used for the subproblems, the condition number of this additive Schwarz method satisfies*

$$\kappa(P) \leq C(1 + \frac{H}{h}),$$

where the constant C is independent of H and h.

4 Implementation

We have implemented this overlapping additive Schwarz solver on two distributed memory architectures of the Intel family, the Caltech Touchstone Delta and the Hypercube iPSC/860 and on the Connection Machine CM5. In this paper, we report on results with the Delta.

A single-program-multiple-data (SPMD) programming model is used. Each (overlapping) subdomain is assigned to one processor. In the preconditioned conjugate gradient method running on each processor, the four major computational kernels are vector updates, inner products, matrix-vector multiply and preconditioning. The vector updates and the inner products are easily parallelized; the latter require communication for the global sum of the local inner products. Since we discretize the problem with a 7 point cell-centered finite difference stencil, the stiffness matrix has only 7 nonzero diagonals and can be stored in local one-dimensional arrays. Therefore the matrix-vector products required in the conjugate gradient iteration are trivially parallelizable. The additive Schwarz preconditioner is naturally parallel. The only nontrivial choice concerns the coarse problem. In our implementation, we use the simplest choice: we construct the coarse matrix in parallel, but we store and solve the coarse problem on all of the processors. For a large number of subdomains, this is not a good choice, since the coarse problem will then dominate the total cost. In that case, we should use a more elaborate implementation of the algorithm and solve the coarse problem in parallel or in a multilevel scheme.

Our fortran implementation is not optimized at any level, since we are interested in the scaling of the algorithm and not in maximizing its performance.

For an analysis of parallelism in domain decomposition, see Gropp [10]. For a general discussion of parallelism in iterative and direct methods, see Demmel, Heath, and Van der Vorst [4].

5 Numerical experiments

We present here the results of some numerical experiments comparing the overlapping additive Schwarz method (ASM) introduced in the previous section and a conjugate gradient method with diagonal preconditioner (JCG). For other experiments with overlapping additive Schwarz methods in three dimensions, see Gropp and Smith [11], Bjørstad, Moe and Skogen [1], Skogen [16] and Cowsar [3].

We discretize the elliptic problem (1) with cell-centered finite differences, which is equivalent to using the lowest order Raviart-Thomas-Nedelec solid rectangular elements with special quadrature rules, see Weiser and Wheeler [17]. We consider

solid rectangular domains Ω. The discrete problem is solved iteratively using a preconditioned conjugate gradient method (ASM or JCG) with zero initial guess. The iteration is stopped when the l^2-norm of the relative residual is less than 10^{-6}. The condition number of the iteration operator is approximated by the Lanczos method. The local and coarse problems in the ASM preconditioner are solved directly with Linpack banded subroutines, which are expensive but very robust. We are in the process of studying the use of approximate local solvers.

Scaling studies in domain decomposition are usually given by keeping the global domain size fixed and increasing the number of processors by decreasing the subdomain size H/h. In this way, the algorithm is changed for every choice of subdomain size and it is difficult to discern among the multiple effects of changes in subdomain size, aspect ratio and surface to volume effects. Moreover, the limitation of the local memory of each processor can severely limit the minimum number of processors used and the global size of the domain.

Another type of scaling consists in keeping the subdomain size H/h fixed and increasing the number of processors by increasing the number of subdomains and therefore increasing the size of the global problem. This is the type of scaling we have chosen in the two following sets of experiments.

(a) "Planar" decomposition with subdomain size $10 \times 10 \times 10$; Ω is decomposed into $(2,2,1),(4,4,1),\cdots,(16,16,1)$ subdomains and the global size varies from $4 \cdot 10^3$ to $256 \cdot 10^3$ unknowns.

(b) "Cubic" decomposition with subdomain size $13 \times 13 \times 13$; Ω is decomposed into $(2,1,1),(4,2,2),\cdots,(10,5,5)$ subdomains and the global size varies from $4,394$ to $549,250$ unknowns.

In both cases, we consider the matrix A diagonal and constant on each subdomain, but varying 4 orders of magnitude across subdomains. More precisely, we "slice" the domain in the x-direction and define $A = I$ on one slice and $diag(A) = (10^3, 10^{-1}, 10)$ in the next. When the matrix A is the identity (Poisson equation), diagonal preconditioning is definitely faster for these problem sizes.

The results for set (a) are reported in Table 1 and 2, and in graphic form in Figure 1, 2, 3, 4. The results for set (b) are reported in Table 3 and 4, and in Figure 5, 6, 7, 8. Each case was run several times to minimize the timing variations due to factors like machine load, numbers of users, etc.. We have observed significant variations (up to 10and we report here the minimum in each case. In the tables and graphs, the total time is the sum of the initialization and iteration times; the last one (iter) is further split into the times for preconditioning (local + coarse), matrix multiplication (A-mult), inner products (lip = local inner product and

psum = *parallel* global sum). We have not timed the vector updates, since they are completely parallel like the local inner products.

It is clear from the results that JCG does not scale well, while ASM scales well until the point where the construction and solution of the coarse problem is dominating. In case (a), this point is around 100 subdomains; in case (b), this point is not yet reached, due to the larger subdomain size. The cause of the bad performance of JCG is that the condition number (and therefore the number of iteration) grows like $1/h^2$, while for ASM it only grows like H/h. Therefore, most of the JCG time is spent iterating doing matrix multiply and global sums for the local inner products, since the preconditioner is completely local (Figure 4 and 8). On the contrary, the initialization time for ASM is considerable, (due to the expensive exact factorization of the local problems and to the parallel construction of the coarse problem, which involves communication) and the iteration time is mostly spent preconditioning (Figure 3 and 7). Note the almost ideal curve for the local part of the ASM preconditioner, while the cost of the coarse problem increases with the number of subdomains. We have run ASM without coarse space in the planar case with 256 subdomains: the condition number increased to 2,737 and the number of iterations to 138, for a total running time of 50.67 sec. (iter = 27.15, local = 23.61, coarse = 0). Therefore the coarse problem is essential if we want a scalable algorithm. However, we again remark that the naive approach consisting in solving the coarse problem on each processor, is not adequate for large problems with many subdomains.

Further studies are being conducted for large heterogeneous problems, the use of inexact solvers for the subdomain problems and nonsymmetric problems.

Acknowledgments. This work was partially supported by the Texas State Grant number 1059 and by the DOE grant DEFG05 - 92ER 25142. We would also like to thank Frédéric d'Hennezel and Lawrence Cowsar for helpful discussions.

N	κ	# it	tot. time	iter.	local	coarse
4=2×2×1	24.94	19	9.13	3.24	2.52	0.38
16=4×4×1	61.34	36	16.27	7.91	6.31	0.99
36=6×6×1	79.59	43	18.08	9.57	7.49	1.32
64=8×8×1	87.92	47	19.84	10.89	8.20	1.83
100=10×10×1	92.15	49	22.15	12.09	8.52	2.63
144=12×12×1	94.54	53	26.38	14.29	9.20	4.06
196=14×14×1	95.99	56	33.49	17.30	9.73	6.36
256=16×16×1	96.89	56	43.71	20.25	9.66	9.42

Table 1: Results for ASM: planar decomposition, $H/h = 10$

N	κ	# it	tot. time	prec.	A-mult	psum	lip
4=2×2×1	216,795	1,238	17.34	1.00	10.70	1.90	1.22
16=4×4×1		2,329	42.37	2.15	25.63	12.77	2.91
36=6×6×1		3,213	60.63	2.97	35.22	19.83	4.09
64=8×8×1		4,122	80.55	3.83	45.20	28.22	5.18
100=10×10×1		4,734	111.13	5.05	56.85	51.03	7.08
144=12×12×1		5,912	135.63	6.14	69.27	60.58	8.43
196=14×14×1		6,846	154.72	6.90	78.38	67.90	9.47
256=16×16×1		7,823	174.63	7.65	88.31	75.33	10.47

Table 2: Results for JCG: planar decomposition, $H/h = 10$

N	κ	# it	tot. time	iter.	local	coarse
2=2×1×1	8.06	12	25.41	5.19	4.36	0.51
16=4×2×2	184.09	57	62.15	32.34	27.71	2.60
54=6×3×3	254.12	73	87.82	49.88	42.99	4.39
128=8×4×4	286.78	81	96.34	57.06	47.75	6.46
250=10×5×5	302.48	83	106.23	62.68	49.00	10.59

Table 3: Results for ASM: cubic decomposition, $H/h = 13$

N	# it	tot. time	prec.	A-mult	psum	lip
2=2×1×1	1,092	25.21	1.62	15.62	1.15	2.13
16=4×2×2	2,323	72.73	3.98	46.80	12.69	5.40
54=6×3×3	3,411	123.83	6.17	77.60	32.50	8.38
128=8×4×4	4,232	171.17	7.63	99.32	60.80	12.13
250=10×5×5	5,137	220.92	9.25	126.48	85.31	14.48

Table 4: Results for JCG: cubic decomposition, $H/h = 13$

Figure 1: Total time

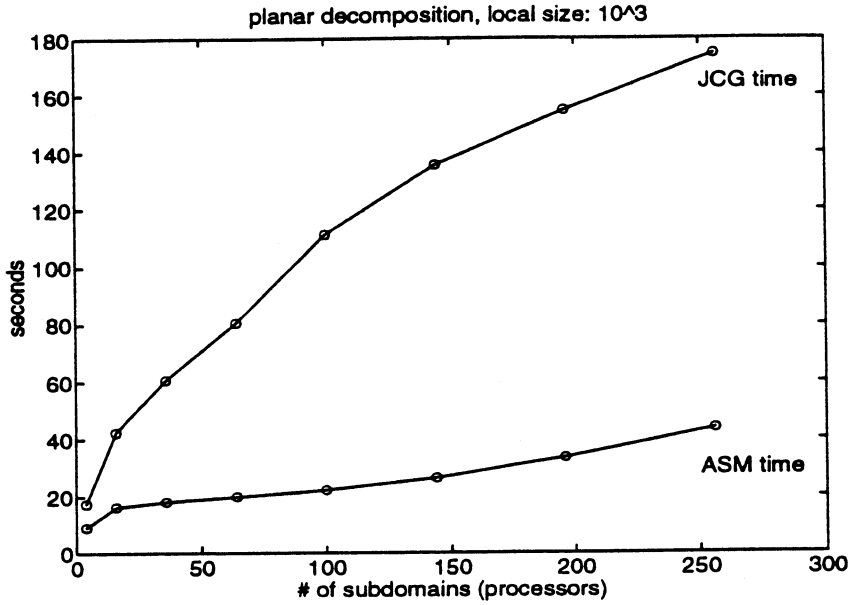

Figure 2: Condition number and iteration count for ASM

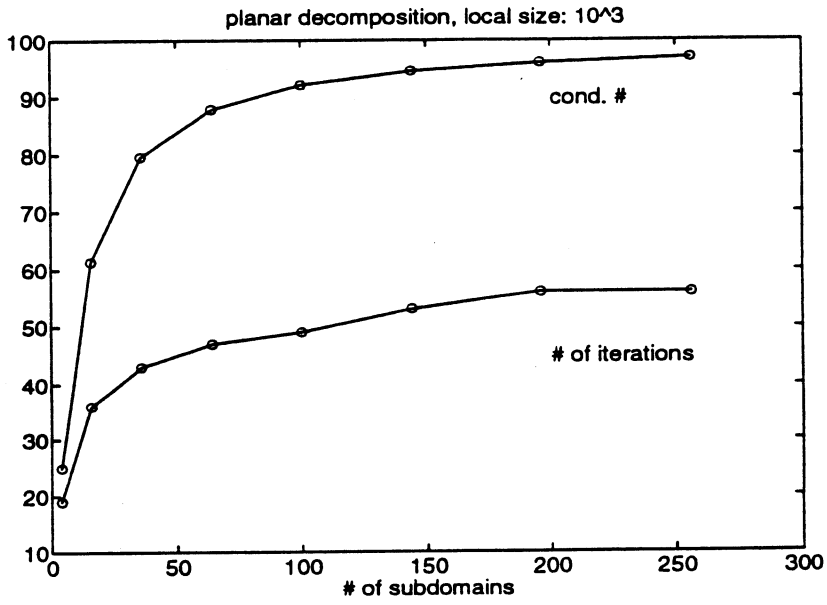

Figure 3: Timings for ASM

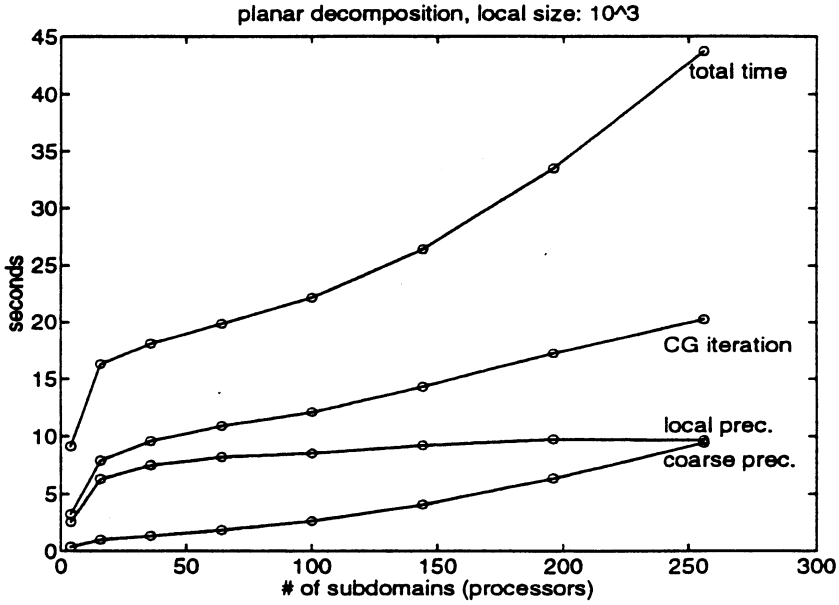

Figure 4: Timings for JCG

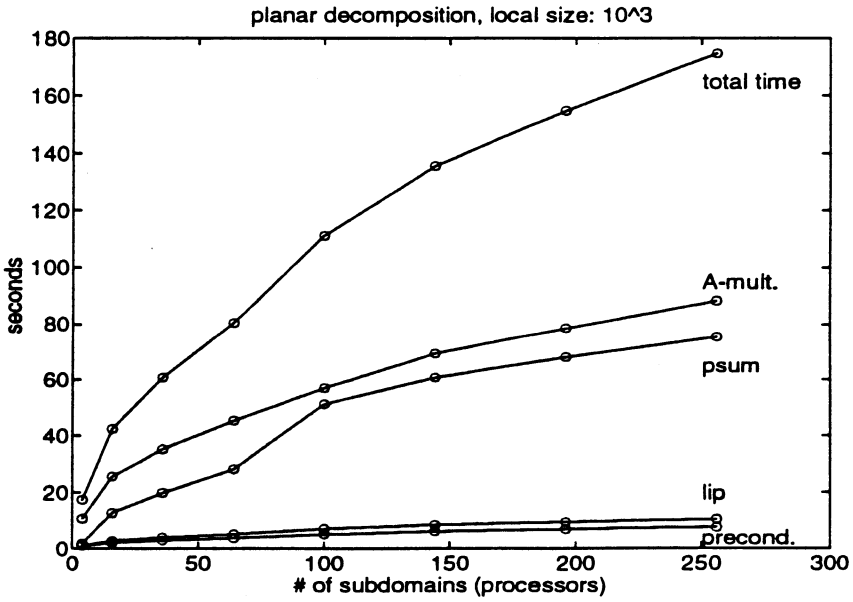

Figure 5: Total time

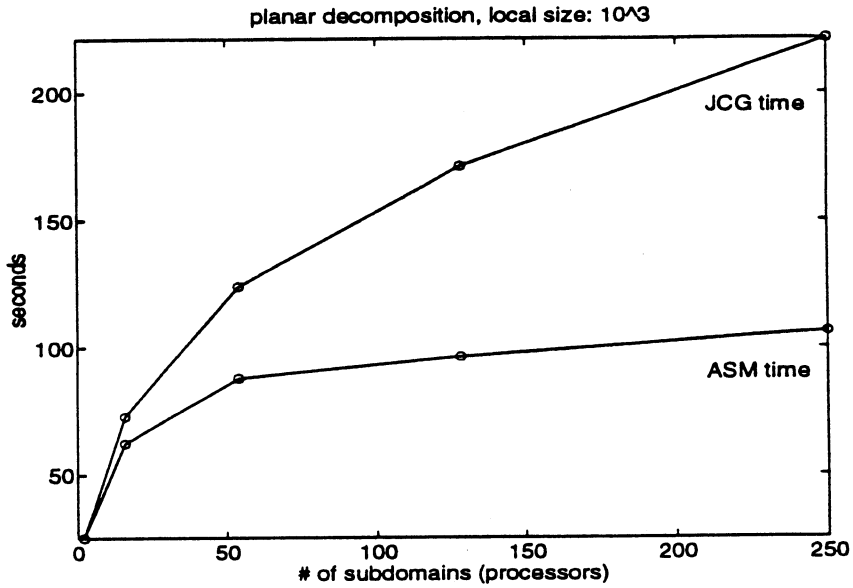

Figure 6: Condition numbers and iteration count for ASM

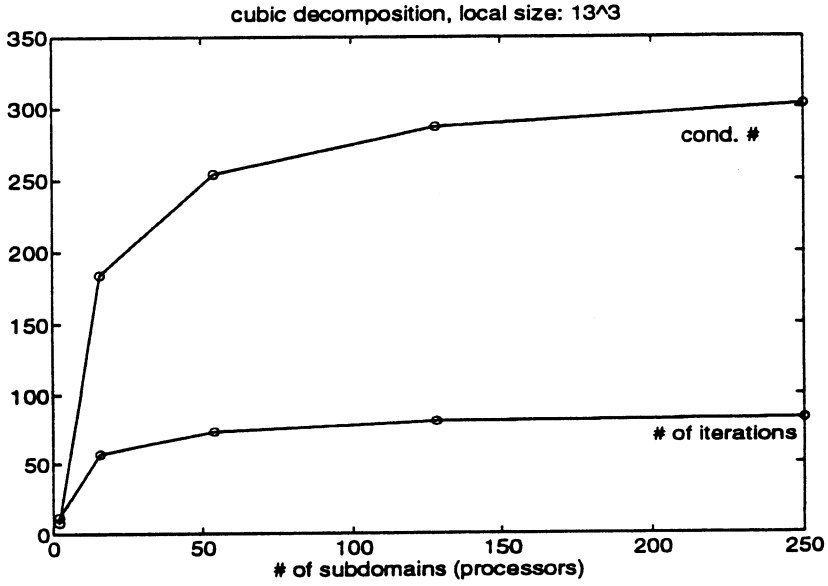

Figure 7: Timings for ASM

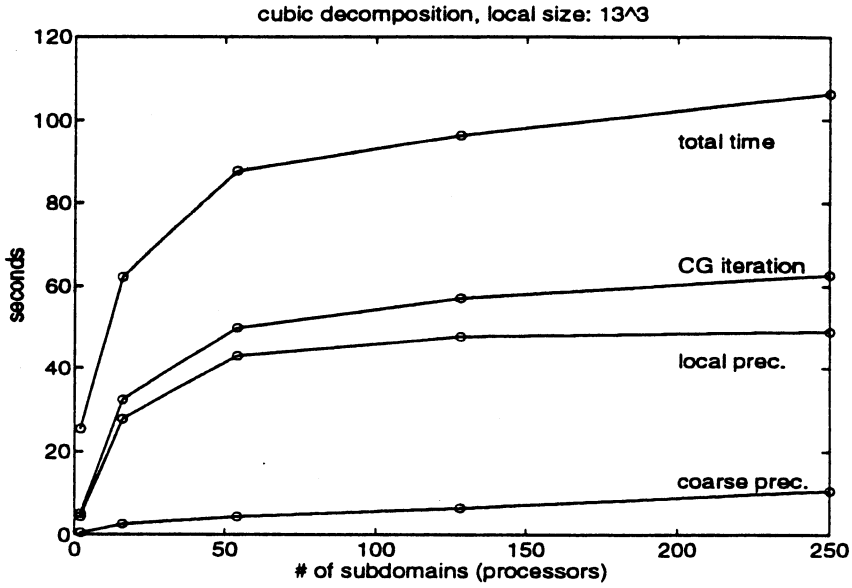

Figure 8: Timings for JCG

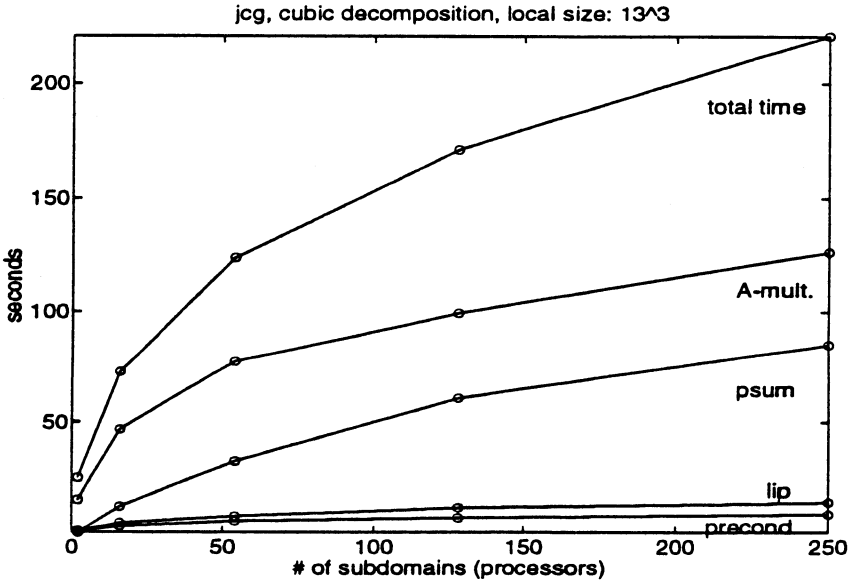

References

[1] P. E. BJØRSTAD, R. MOE, AND M. SKOGEN, *Parallel domain decomposition and iterative refinement algorithms*, in Parallel Algorithms for PDEs, Proceedings of the 6th GAMM-Seminar held in Kiel, Germany, January 19–21, 1990, W. Hackbusch, ed., Braunschweig, Wiesbaden, 1990, Vieweg-Verlag.

[2] F. BREZZI, *On the existence, uniqueness and approximation of saddle-point problems arising from lagrangian multipliers*, RAIRO Anal. Numer., 8 (1974), pp. 129–151.

[3] L. C. COWSAR, *Dual variable schwarz methods for mixed finite elements*, Tech. Rep. TR9309, Rice University, Department of Computational and Applied Mathematics, March 1993.

[4] J. W. DEMMEL, M. T. HEATH, AND H. A. V. D. VORST, *Parallel numerical linear algebra*, Acta Numerica 1993, (1993), pp. 111–197.

[5] M. DRYJA, B. F. SMITH, AND O. B. WIDLUND, *Schwarz analysis of iterative substructuring algorithms for elliptic problems in three dimensions*, tech. rep., Department of Computer Science, Courant Institute, 1993.

[6] M. DRYJA AND O. B. WIDLUND, *Domain decomposition algorithms with small overlap*, Tech. Rep. 606, Department of Computer Science, Courant Institute, May 1992. To appear in SIAM J. Sci. Stat. Comput.

[7] ——, *Schwarz methods of Neumann-Neumann type for three-dimensional elliptic finite element problems*, Tech. Rep. 626, Department of Computer Science, Courant Institute, March 1993. Submitted to Comm. Pure and Appl. Math.

[8] R. E. EWING AND J. WANG, *Analysis of the Schwarz algorithm for mixed finite element methods*, RAIRO Model. Math. Anal. Numer., 26 (1992), pp. 739–756.

[9] R. FREUND, G. H. GOLUB, AND N. NACHTIGAL, *Iterative solution of linear systems*, Acta Numerica 1992, (1992), pp. 57–100.

[10] W. D. GROPP, *Parallel computing and domain decomposition*, in Fifth Conference on Domain Decomposition Methods for Partial Differential Equations, T. F. Chan, D. E. Keyes, G. A. Meurant, J. S. Scroggs, and R. G. Voigt, eds., Philadelphia, PA, 1992, SIAM.

[11] W. D. GROPP AND B. F. SMITH, *Experiences with domain decomposition in three dimensions: Overlapping Schwarz methods*, tech. rep., Mathematics and Computer Science Division, Argonne National Laboratory, 1992. To appear in the Proceedings of the Sixth International Symposium on Domain Decomposition Methods.

[12] T. P. MATHEW, *Domain Decomposition and Iterative Refinement Methods for Mixed Finite Element Discretisations of Elliptic Problems*, PhD thesis, Courant Institute of Mathematical Sciences, September 1989. Tech. Rep. 463, Department of Computer Science, Courant Institute.

[13] ——, *Schwarz alternating and iterative refinement methods for mixed formulations of elliptic problems, part I: Algorithms and Numerical results*, Numerische Mathematik, (1993). To appear.

[14] ——, *Schwarz alternating and iterative refinement methods for mixed formulations of elliptic problems, part II: Theory*, Numerische Mathematik, (1993). To appear.

[15] J. E. ROBERTS AND J.-M. THOMAS, *Mixed and hybrid methods*, in Handbook of Numerical Analysis, vol. 2, P. G. Ciarlet and J. L. Lions, eds., North-Holland, 1991.

[16] M. D. SKOGEN, *Schwarz Methods and Parallelism*, PhD thesis, Department of Informatics, University of Bergen, Norway, February 1992.

[17] A. WEISER AND M. F. WHEELER, *On convergence of block-centered finite difference for elliptic problems*, SIAM J. Numer. Anal., 25 (1988), pp. 351–357.