Computer Science 119

Lubin Vulkov, Jerzy Waśniewski
Plamen Yalamov (Eds.)

Numerical Analysis and Its Applications

First International Workshop, WNAA'96
Rousse, Bulgaria, June 1996

Springer

High Performance Solution of Linear Systems Arising from Conforming Spectral Approximations for Non-Conforming Domain Decompositions

Marcin Paprzycki[1] and Andreas Karageorghis[2]

Department of Mathematics and Statistics, University of Cyprus
1678 Nicosia, Cyprus
andreask@pythagoras.mas.ucy.ac.cy

Abstract. We apply a conforming spectral collocation technique to non-conforming domain decompositions. The resulting global matrices have a particular block structure. We study the performance of various direct methods of solution of the resulting linear system on a RS6000 workstation, a SGI Power Challenge and a Cray J-916 supercomputer.

1 Introduction

In this paper we study the performance of three direct methods for the solu of the global systems resulting from spectral approximations for a certain (of domain decompositions. In particular, we examine the linear systems ari from conforming spectral approximations in non-conforming decompositio rectangular domains, developed in [7]. The spectral approximations which used are conforming, that is, the solution is C^0 continuous at all points a the subdomain interfaces for second order problems and C^1 continuous a points across the subdomain interfaces for fourth order problems. The mat resulting from these approximations possess a specific block diagonal struc which can be exploited by applying a banded system solver [8] or a capacita type technique [5, 6]. The performance of these two approaches is compar the performance of two full matrix solvers from the NAG library [8].

2 Domain decomposition and spectral approximation

We consider the problem

$$\nabla^2\phi(x,y) = F(x,y)$$

on the rectangle $(\alpha,\beta)\times(a,b)$ subject to Dirichlet boundary condition shall assume that the boundary conditions can be expressed as polynomial in [5, 7], for the partitions $\alpha = \alpha_0 < \alpha_1 < \alpha_2 < \ldots < \alpha_{N-1} < \alpha_N =$ $a = a_0 < a_1 < a_2 < \ldots < a_{N-1} < a_N = b$, $N \in \mathbb{N}$, we consider the foll decomposition: the rectangle $(\alpha,\beta)\times(a,b)$ is decomposed into $2N-1$ subdo in the following way: for $k = 1,2,\ldots,N-1$, subdomain $2k-1$ is the rect

$_{k-1}, \alpha_k) \times (a_{k-1}, a_N)$ and subdomain $2k$ is the rectangle $(\alpha_k, \alpha_N) \times (a_{k-1}, a_k)$. ubdomain $2N-1$ is the rectangle $(\alpha_{N-1}, \alpha_N) \times (a_{N-1}, a_N)$. In each subdomain he solution is approximated by

$$\phi_s(x, y) = \sum_{m=0}^{M_s} \sum_{n=0}^{N_s} \gamma^s_{mn} \hat{T}^s_m(x) \tilde{T}^s_n(y), \quad s = 1, 2, ..., 2N-1, \qquad (2.2)$$

here the functions $\hat{T}^s_m(x)$ and $\tilde{T}^s_n(y)$ are the shifted Chebyshev polynomials de- ned on the corresponding intervals of each region and the collocation points n each interval of each region (e.g. $\{x^s_i\}^{M_s}_{i=0}$) are the Gauss-Lobatto points 1, 2]. We shall assume that $M_{2k} \leq \min\{M_{2k+1}, M_{2k+2}\}$ and that $N_{2k-1} \leq$ in$\{N_{2k}, N_{2k+1}\}$, $k = 1, 2, ..., N-1$

For for the above problem and domain decomposition it can be shown that e spectral approximation (2.2) yields C^0 conforming approximations on all the bdomain interfaces [5, 7].

Methods of solution

.1 Capacitance technique

he structure of the global system for the five subdomain decomposition is of e form given in Figure 1.

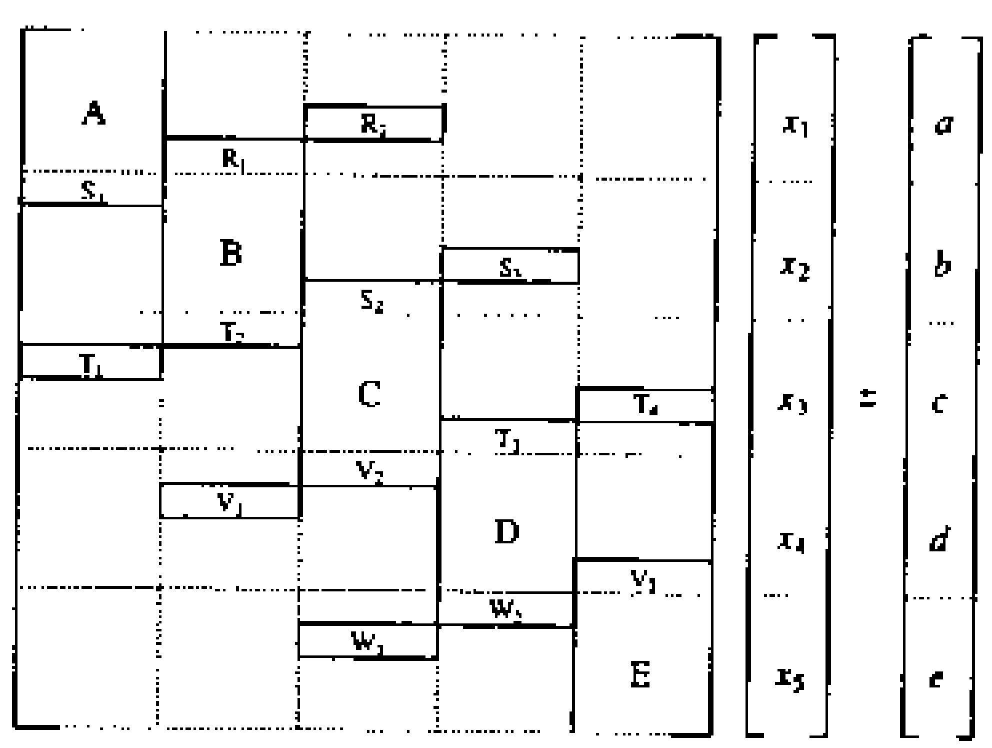

Fig. 1. Matrix for the 5-subdomain decomposition.

From the above system (Figure 1) we may express $\underline{x}_1$ and $\underline{x}_5$ in terms of $\underline{x}_2$, $\underline{x}_3$ and $\underline{x}_4$. We then substitute these expressions into the original system to yield a reduced system of the form:

$$\begin{aligned} \bar{A}_2\underline{x}_2 + \bar{S}_2\underline{x}_3 + \bar{S}_3\underline{x}_4 &= \bar{\alpha}_2 \\ \bar{T}_2\underline{x}_2 + \bar{A}_3\underline{x}_3 + \bar{T}_3\underline{x}_4 &= \bar{\alpha}_3 \\ \bar{V}_1\underline{x}_2 + \bar{V}_2\underline{x}_3 + \bar{A}_4\underline{\bar{x}}_4 &= \bar{\alpha}_4 \end{aligned}$$

The above system is then solved for $\underline{x}_2, \underline{x}_3$ and $\underline{x}_4$. The unknown vectors $\underline{x}_1$ and $\underline{x}_5$ may be obtained by back subsitution. This process can be easily generalized for any non-conforming multidomain decomposition [6].

3.2 Banded solvers

As can be seen from Figure 1 the linear system can be viewed as banded with bandwith equal to $\max\{(\sum_{n=0}^{2}(M_{S-2+K}-1)(N_{S-2+K}+1)), (\sum_{n=0}^{2}(M_{S+K}+1)(N_{S+K}+1))\}$ (the bandwidth is independent of the number of subdomains in the decomposition). It can be also observed that the banded approach may not be the most efficient as there is a substantial number of zeros inside the band that the solver cannot exploit.

4 Experimental results

4.1 Numerical example

The performance of the various techniques described in Section 3 was tested on the following test problem:

$$\nabla^2\phi(x, y) = (y^2-1)e^x + (x^2-1)e^y + 2e^x + 2e^y \quad \text{on} \quad (-1, 1)^2$$

subject to Dirichlet boundary conditions which correspond to the exact solution of this problem $\phi(x, y) = (y^2-1)e^x + (x^2-1)e^y$. We used the decomposition (in the notation of Section 2), $\alpha_i = \alpha_{i-1} + (1/2)(\alpha_N - \alpha_{i-1}), i = 1, 2, ..., N-1, \alpha_0 = \alpha = -1, \alpha_N = \beta = 1$ and $a_i = a_{i-1} + (1/2)(a_N - a_{i-1}), i = 1, 2, ..., N-1, a_0 = a = -1, a_N = b = 1$. We also took (in Equation (2.2)) $M_s = N_s = n, s = 1, 2, ..., L$. The total number of unknowns (u_T) is therefore $u_T = L(n+1)^2$.

4.2 Implementation

The experiments were performed on a RS6000-550 workstation, a SGI Power Challenge and a Cray J-916 supercomputer. Timings on the RS6000 were obtained on an empty machine using the *time* function, on the Power Challenge the timings were obtained using the system timer *getrusage* and on the Cray using the *perftrace* utility. Each result presented here is an average of multiple runs. We experimented with two dense solvers. The first is the pre-packaged general solver routine F04ATF from the NAG library [8], which performs an LU decomposition, solves the linear system and if necessary improves the solution by iterative refinement(s). The second is a LAPACK [4] based pair of routines F07ADF (LAPACK routine _GETRF) performing the LU decomposition and F07AEF (LAPACK routine _GETRS) solving the factorized system. We used both these approaches inside the capacitance technique. For the banded solver we used the F07BDF/F07BEF pair (LAPACK pair _GBTRF/ _GBTRS).

4.3 Dense Solver Performance

In Table 1 the timings in seconds for $n = 4, 8, 12$ are presented for the dense solvers for the three computers we have experimented with for the three, five, seven and nine element decompositions.

		RS6000		*SGI*		*Cray*	
n	u_T	*F04*	*F07*	*F04*	*F07*	*F04*	*F07*
			Three elements				
4	75	0.08	0.06	0.02	0.02	0.04	0.06
8	243	0.80	0.53	0.19	0.08	0.33	0.24
12	507	7.47	4.83	0.94	0.42	1.40	0.87
			Five elements				
4	125	0.15	0.11	0.05	0.03	0.11	0.11
8	405	3.11	1.63	0.53	0.24	0.89	0.56
12	845	22.83	15.20	3.43	1.19	4.48	2.81
			Seven elements				
4	175	0.30	0.18	0.09	0.05	0.19	0.16
8	567	7.31	3.38	1.18	0.55	1.79	1.11
12	1183	49.35	31.09	8.94	5.51	10.20	6.76
			Nine elements				
4	225	0.51	0.23	0.16	0.08	0.30	0.23
8	729	11.41	5.93	2.25	1.10	3.18	1.96
12	1521	77.75	51.64	18.72	12.50	19.30	13.60

Table 1. Timings for the full matrix solution

It can be observed that the LAPACK pair outperforms the F04ATF solver for all machines and all decompsition sizes. This can be related to the iterative refinements which when performed are a series of level 2 BLAS operations which are characterized by a much poorer performance than the level 3 BLAS based LAPACK kernels [4, 9].

It can also be observed that as the matrix size u_T increases the Cray's performance reaches that of the Power Challenge. This fact is rather peculiar as one processor of the Cray J-916 reaches only 195 MFlops in comparison with 270 MFlops of the practical peak performance of the Power Challenge. There are two points that need to be raised: Firstly, the performance of the SGI has been observed to decrease as the matrix size increases, whereas the performance of the Cray remains unchanged. Secondly, the BLAS/LAPACK kernels are much better optimized by the Cray then the same routines in the SGI's scientific computing library. The crossing point where the Cray outperforms the Power Challenge occurs at $n = 14$ (for the nine element decomposition).

Finally, for the largest systems the Cray's performance reaches 187 MFlops, which is approximately 95% of the practical peak.

4.4 Banded Solver Performance

It is clear that for the three element decomposition the banded structure is non existent. This is why we did not applied the banded solver in that case. In Table 2 the results for the banded technique are summarized for the five and seven element decompositions for $n = 4, 7, 10, 13, 16$ (the results of the nine element decompostion are presented in section 4.5).

n	*Bandwidth*	*RS6000*	*SGI*	*Cray*
		Five	*elements*	
4	75	0.09	0.03	0.11
7	192	0.94	0.12	0.32
10	363	4.96	0.52	0.95
13	507	19.06	2.08	2.83
16	867		7.92	7.85
		Seven	*elements*	
4	75	0.16	0.04	0.15
7	192	1.81	0.25	0.56
10	363	10.15	1.24	2.03
13	507	38.84	5.72	6.78
16	867		20.06	19.10

Table 2. Timings for the banded solver

The missing result for the RS6000 means that the system of this size did not fit into the memory. When compared with the results from Table 1 it can be observed that the performance of the banded solver is approximately 1.5 times better. This is most apparent for the RS6000 and least visible for the SGI. As above, the increase in matrix size improves the performance of the Cray relatively to that of the SGI.

4.5 Capacitance Technique Performance

As was the case in the banded system solver approach, the capacitance technique as described in Section 3 cannot be applied to the three-domain decomposition. Based on the results of Section 4.3 we applied the capacitance technique with the F07ADF/F07AEF pair. Further experiments with F04ATF confirmed the validity of this choice. The performance gain in the largest cases was about 10%. In Table 3 the results for the banded technique are summarized for the five and seven element decompositions for $n = 4, 7, 10, 13, 16$ (the nine element decompostion results are presented below).

n	u_T	*RS6000*	*SGI*	*Cray*
		Five	*elements*	
4	125	0.07	0.02	0.07
7	320	0.46	0.09	0.21
10	605	2.35	0.30	0.55
13	980	9.67	0.99	1.49
16	2601		2.94	3.78
		Seven	*elements*	
4	125	0.13	0.03	0.08
7	320	0.69	0.13	0.26
10	605	3.95	0.64	0.75
13	980	17.17	2.19	2.24
16	2601		6.86	5.98

Table 3. Timings for the capacitance technique

Finally, in Table 4 we summarize the nine element performance for the banded solver and the capacitance technique for all three computers and for $n = 4, 5, \ldots, 16$.

	RS6000		*SGI*		*Cray*	
n	*banded*	*capacit*	*banded*	*capacit*	*banded*	*capacit*
4	0.28	0.10	0.07	0.05	0.20	0.11
5	0.70	0.23	0.13	0.08	0.33	0.18
6	1.49	0.53	0.24	0.15	0.53	0.31
7	3.09	1.19	0.45	0.28	0.86	0.55
8	5.31	2.18	0.90	0.51	1.38	0.88
9	9,79	3.95	1.56	0.88	2.27	1.33
10	16.64	6.93	2.61	1.48	3.61	1.93
11	26.45	11.54	4.57	2.43	5.70	2.82
12	42.68	18.08	8.20	3.81	8.57	3.63
13	64.09	29.83	13.24	5.82	12.80	5.91
14			20.82	8.58	18.90	7.23
15			31.41	12.67	29.00	26.00
16			55.71	18.18	38.93	13.10

Table 4. Timings for the capacitance and banded solvers

Comparing results in Tables 2, 3 and 4 one can observe that the capacitance echnique outperforms the banded solver in all cases. In the largest case (nine ubdomain decomposition) the performance gain is about 2.15 for the RS6000, .06 for the SGI Power Challenge and 2.97 for the Cray. The results reconfirm he fact that as the matrix size increases the Cray performs better than the

SGI Power Challenge. For the largest reported case the performance gain is 1.43 for the banded solver and 1.39 for the capacitance technique. In this case the Cray's banded solver reaches 163 MFlops and the capacitance technique runs at 112 MFlops (83% and 57% of the practical peak perfomance respectively). The Cray's performance drop for $n = 15$ is related to the memory bank conflicts (the blocks are of sizes that are multiples of 16; see [4, 9]).

5 Conclusions

In this study we present efficient direct methods of solution for the global systems resulting from conforming approximations for a class of non-conforming domain decompositions. A comparison of these methods is carried out for several decompositions and the results indicate that a capacitance-type technique which expoits the block structure is best suited for the solution of such systems. Our results also indicate that a substantial improvement is necessary in the SGI's scientific computing library for it to be able to compete with the Cray in terms of performance. The next step of our research is to study the performance characteristics of the general sparse multifrontal solver UMFPACK [3] when applied to the problem in question.

6 Acknowledgement

This work was partially supported by a grant from the University of Cyprus. The computer time grant from the NCSA at Urbana-Champain is kindly acknowledged.

References

1. Boyd J. P.: Chebyshev and Fourier Spectral Methods. Springer-Verlag, New York (1989)
2. Canuto C., Hussaini M., Quarteroni A., Zang T.: Spectral Methods in Fluid Dynamics. Springer-Verlag, New York (1988)
3. Davis T. A., Duff I. S.: A combined unifrontal/multifrontal method for unsymmetric sparse matrices. Technical Report TR-95-020, CISE Department, University of Florida (1995)
4. Dongarra J. J., Duff J. S., Sorensen D. C., van der Vorst H. A.: Solving Linear Systems on Vector and Shared Memory Computers. SIAM, Philadelphia (1991)
5. Karageorghis A.: Conforming Spectral Domain Decomposition Schemes. (this volume)
6. Karageorghis A., Paprzycki M.: An efficient direct method for fully conforming spectral collocation schemes. Annals of Numerical Mathematics. (to appear)
7. Karageorghis A., Sivaloganathan S.: Conforming spectral approximations for non-conforming domain decompositions. Technical Report TR/07/96, Department of Mathematics and Statistics, University of Cyprus (1996)
8. Numerical Algorithms Group Library Mark 15. NAG(UK) Ltd, Wilkinson House Jordan Hill Road, Oxford, UK
9. Paprzycki M., Cyphers C.: Multiplying matrices on the Cray – Practical Considerations, CHPC Newsletter 6 (1991) 77-82