



HAL
open science

Parallel GMRES with a multiplicative Schwarz preconditioner

Désiré Nuentsa Wakam, Guy-Antoine Atenekeng Kahou

► **To cite this version:**

Désiré Nuentsa Wakam, Guy-Antoine Atenekeng Kahou. Parallel GMRES with a multiplicative Schwarz preconditioner. [Research Report] RR-7342, 2010. inria-00508277v1

HAL Id: inria-00508277

<https://inria.hal.science/inria-00508277v1>

Submitted on 2 Aug 2010 (v1), last revised 23 Sep 2010 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

***Parallel GMRES
with a multiplicative Schwarz preconditioner***

Desire NUENTSA WAKAM – Guy Antoine ATENEKENG KAHOU

N° 7342

Juillet 2010

Observation and Modeling for Environmental Sciences



R *apport
de recherche*

Parallel GMRES with a multiplicative Schwarz preconditioner

Desire NUENTSA WAKAM – * Guy Antoine ATENEKENG KAHOU†

Theme : Observation and Modeling for Environmental Sciences
Équipes-Projets Sage

Rapport de recherche n° 7342 — Juillet 2010 — 24 pages

Abstract: In this paper, we present an hybrid solver for linear systems that combines a Krylov subspace method as accelerator with some overlapping domain decomposition method as preconditioner. The preconditioner uses an explicit formulation associated to one iteration of the classical multiplicative Schwarz method. To avoid communications and synchronizations between subdomains, the Newton-basis GMRES implementation is used as accelerator. This requires to divide the computation of the orthonormal basis in two steps: the preconditioned Newton basis is computed then it is orthogonalized. The first step is merely a sequence of matrix-vectors and solution of linear systems associated to subdomains. We describe the fine-grained parallelism that is used in these kernel operations. The second step uses a parallel implementation of dense QR factorization on the resulted basis.

At each application of the preconditioner operator, local systems associated to the subdomains are solved with some accuracy depending on the global physical problem. We show that this step can be further parallelized with calls to external third-party solvers. To this end, we define two levels of parallelism in the solver: the first level is intended for the computation and the communication across all the subdomains; the second level of parallelism is used inside each subdomain to solve the smaller linear systems induced by the preconditioner.

Numerical experiments are performed on several problems to demonstrate the benefits of such approaches, mainly in terms of global efficiency and numerical robustness.

Key-words: domain decomposition, preconditioning, multiplicative Schwarz, Parallel GMRES, Newton basis, multilevel parallelism

* INRIA, Campus universitaire de Beaulieu, 35042 Rennes Cedex. E-mail : desire.nuentsa_wakam@irisa.fr

† INRIA and LRI, Parc Orsay Université, 91893 Orsay Cedex. E-mail : atenekeng@yahoo.com

GMRES parallèle préconditionné par Schwarz multiplicatif

Résumé : Dans cet article, nous présentons un solveur hybride pour la résolution des systèmes linéaires sur des architectures parallèles. Le solveur associe un accélérateur basée sur les sous-espaces de Krylov à un préconditionneur basé sur une décomposition de domaine. Le préconditionneur est défini à partir d'une formulation explicite correspondant à une itération de Schwarz multiplicatif. L'accélérateur est de type GMRES. Dans le but de réduire la communication entre les sous-domaines, nous utilisons une version parallèle de GMRES qui divise en deux étapes la construction de la base du sous-espace de Krylov associée: La première étape consiste à générer les vecteurs de la base dans un pipeline d'opérations à travers tous les processeurs. Les principales opérations ici sont les produits matrice-vecteur et les résolutions des sous-systèmes linéaires dans les sous-domaines. La deuxième étape permet d'effectuer une factorisation QR parallèle sur une matrice rectangulaire formée par les vecteurs construits précédemment; ceci permet d'obtenir une base orthogonale du sous-espace de Krylov.

Les sous-systèmes issus de l'application du préconditionneur sont résolus à différents niveaux de précisions par une factorisation LU ou ILU, en fonction de la difficulté du problème sous-jacent. De plus, cette étape peut être résolue en parallèle via des appels aux solveurs parallèle externes. Pour cela, nous utilisons deux niveaux de parallélisme dans notre solveur global. Le premier niveau permet de définir les calculs et les communications à travers tous les sous-domaines. Le deuxième niveau est définie à l'intérieur de chaque sous-domaine pour la résolution des sous-systèmes issus du préconditionneur.

Plusieurs tests numériques sont effectués pour valider l'efficacité de cette approche, en particulier pour les aspects d'efficacité et de robustesse.

Mots-clés : Décomposition de domaine, preconditionnement, Schwarz multiplicatif, GMRES parallèle, Base de Newton, parallélisme multiniveaux.

1 Introduction

In this paper, we are interested in the parallel computation of the solution of the linear system (1)

$$Ax = b \tag{1}$$

with $A \in \mathbb{R}^{n \times n}$, $x, b \in \mathbb{R}^n$. Over the two past decades, the GMRES iterative method proposed by Saad and Schultz [33] has been proved very successful for this type of systems, particularly when A is a large sparse nonsymmetric matrix. Usually, to be robust, the method solves a preconditioned system (2)

$$M^{-1}Ax = M^{-1}b \quad \text{or} \quad AM^{-1}(Mx) = b \tag{2}$$

where M^{-1} is a preconditioner operator that accelerate the convergence of the iterative method.

On computing environments with a distributed architecture, preconditioners based on domain decomposition are of natural use. Their formulation reduces the global problem to several subproblems, where each subproblem is associated to a subdomain; therefore, one or more subdomains are associated to a node of the parallel computer and the global system is solved by exchanging information between neighboring subdomains. Generally, in domain decomposition methods, there are two ways of deriving the subdomains : (i) from the underlying physical domain and (ii) from the adjacency graph of the coefficient matrix A . In any of these partitioning techniques, subdomains may overlap. Overlapping domain decomposition approaches are known as Schwarz methods while non-overlapping approaches refer to Schur complement techniques. Here, we are interested in preconditioners based on the first class. Depending on how the global solution is obtained, the Schwarz method is *additive* or *multiplicative* [38, Ch. 1]. The former approach computes the solution of subproblems simultaneously in all subdomains. It is akin to the block Jacobi method; therefore, the additive Schwarz method has a straightforward implementation in a parallel environment [13]. Furthermore, it is often used in conjunction with Schur complement techniques to produce hybrid preconditioners [14, 23, 34].

The multiplicative Schwarz method builds a solution of the global system by alternating successively through all the subdomains; it is therefore similar to the block Gauss-Seidel method on an extended system; Thus, compared to the additive approach, it will theoretically require fewer iterations to converge. However, good efficiency is difficult to obtain in a parallel environment due to the high data dependencies between the subdomains. The traditional approach to overcome this is through graph coloring by associating different colors to neighboring subdomains. Hence, the solution in subdomains of the same color could be updated in parallel [38, Ch. 1]. Recently, a different approach has been proposed [4, 6]. In that work, the authors proposed an explicit formulation associated to one iteration of the multiplicative Schwarz method. This formulation requires that the matrix is partitioned in block diagonal form [5] such that each block has a maximum of two neighbors; from this explicit formula, the residual vector is determined out of the computation of the new approximate global solution, and therefore could be parallelized through sequences of matrix-vectors and local solutions in subdomains.

The first purpose of this paper is to present the parallelism that is obtained when this explicit formulation is used to build a preconditioned Krylov basis for the GMRES method. This is achieved through the Newton basis implementation proposed in [7] and applied in [18, 37]. Hence, the usual inner products and global synchronizations are avoided across all the subdomains and the resulted algorithm leads to a pipeline parallelism. We will refer to this as a *first-level of parallelism* in GMRES. The second and main purpose of our work here

is to further use parallel operations when solving local systems associated to subdomains. To this end, we introduce a *second-level of parallelism* to deal with those subsystems. The first observation is that for systems that arise from non-elliptic operators, the number of subdomains should be small to guarantee the convergence. Consequently, the local systems could be very large and should be solved efficiently to accelerate the convergence. The second observation comes from the architecture of the current parallel computers. These two levels of parallelism enable an efficient usage of the allocated resources by dividing tasks across and inside all the available nodes. This approach has been recently used to enhance scalability of hybrid preconditioners based on additive Schwarz and Schur Complement techniques [22].

The remaining part of this paper is organized as follows. Section 2 recall the explicit formulation of the Multiplicative Schwarz preconditioner. After that, a parallel implementation of the preconditioned Newton-basis GMRES is given. Section 3 provides the second level of parallelism introduced to solve linear systems in subdomains. As those systems should be solved several times with different right hand sides, the natural way is to use a parallel third-party solver based on an (incomplete) LU factorization. The parallelism in these solvers are intended for distributed-memory computes; we will discuss the benefits of using such approach. In section 4, we provide intensive numerical results that reveal good performance of this parallel hybrid solver. The matrices of tests are taken either from public academic repositories or from industrial test cases. Concluding remarks and future directions of this work are given at the end of the document.

2 A parallel version of GMRES preconditioned by multiplicative Schwarz

In this section, the explicit formulation of the multiplicative Schwarz method is introduced. Then we show how to use it efficiently as a parallel preconditioner for the GMRES method. A good parallelism is obtained thanks to the use of the Newton-Krylov basis.

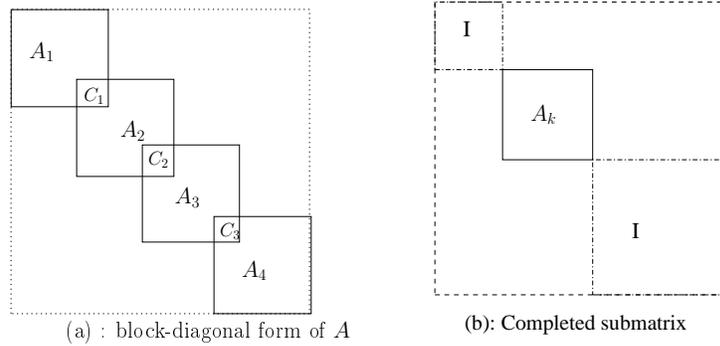
2.1 Explicit formulation of the multiplicative Schwarz preconditioner

From the equation (1), we consider a permutation of the matrix A into p overlapping partitions A_i . We denote by C_i the overlapping matrix between A_i and A_{i+1} ($i = 1, \dots, p-1$). Here, each diagonal block has a maximum of two neighbors (see Figure 1.a). Assuming that there is no full line (or column) in the sparse matrix A , such partitioning can be obtained from the adjacency graph of A by means of profile reduction and level sets [5]. We define \bar{A}_i (resp. \bar{C}_i) the matrix A_i (resp. C_i) completed by identity to the size of A (see Figure 1.b).

If \bar{A}_i and \bar{C}_i are nonsingular matrices, then the matrix associated to one iteration of the classical multiplicative Schwarz method is defined [6] as :

$$M^{-1} = \bar{A}_p^{-1} \bar{C}_{p-1} \bar{A}_{p-1}^{-1} \bar{C}_{p-2} \dots \bar{A}_2^{-1} \bar{C}_1 \bar{A}_1^{-1} \quad (3)$$

This explicit formulation is very useful to provide stand-alone algorithm for the essential operation $y \leftarrow M^{-1}x$ used in iterative methods. However, since dependencies between subdomains are still present in this expression, no efficient parallel algorithm could be obtained for this single operation. Now, if a sequence of vectors v_i should be generated such that $v_i \leftarrow M^{-1}Av_{i-1}$, then the algorithm would produce a very good pipeline parallelism. For the preconditioned Krylov method, these vectors are simply those that span the Krylov subspace. However, to get a parallel process, the basis of this Krylov subspace is formed with the Newton polynomials as described in the next section.


 Figure 1: Partitioning of A in four subdomains

2.2 Background on GMRES with the Newton basis

A preconditioned restarted GMRES(m) method minimizes the residual vector $r_m = M^{-1}(b - Ax_m)$ in the Krylov subspace $x_0 + \mathcal{K}_m$ where x_0 is the initial approximation, x_m the current iterate and $\mathcal{K}_m = \text{span}\{r_0, M^{-1}Ar_0, \dots, (M^{-1}A)^m r_0\}$. The new approximation is of the form $x_m = x_0 + V_m y_m$ where y_m minimizes r_m . The most time consuming part in this method is the construction of the orthonormal basis V_m of \mathcal{K}_m ; the Arnoldi process is generally used for this purpose [32]. It constructs the basis and orthogonalizes it in the same time. The effect of this is the presence of global communication between all the processes. Hence, no parallelism could be obtained across the subdomains with this approach. However, synchronization points can be avoided by decoupling the construction of V_m into two independent phases : first the basis is generated *a priori* then it is orthogonalized.

Many authors proposed different ways to generate this *a priori* basis [7, 40, 16]. For stability purposes, we choose the Newton basis introduced by Bai et al [7] and used by Erhel [18]:

$$\widehat{V}_{m+1} = [\mu_0 r_0, \mu_1 (M^{-1}A - \lambda_1 I)r_0, \dots, \mu_m \prod_{j=1}^m (M^{-1}A - \lambda_j I)r_0] \quad (4)$$

where μ_j and λ_j ($j = 1, \dots, m$) are respectively scaling factors and approximate eigenvalues of A . With a chosen initial vector $v_0 = r_0 / \|r_0\|$, a sequence of vectors v_1, v_2, \dots, v_m is first generated such that:

$$v_j = \sigma_j (M^{-1}A - \lambda_1 I)v_{j-1} \quad j = 1 \dots m \quad (5)$$

where

$$\sigma_j = 1 / \|(M^{-1}A - \lambda_j I)v_{j-1}\| \quad (6)$$

To avoid global communication, the vectors v_j are normalized at the end of the process. Hence, the scalars μ_j from the equation (4) are easily computed as the product of scalars σ_j . These steps are explained in detail in the section 2.2.1 and 2.2.2. At this point, we get a normalized basis V_m such that

$$M^{-1}AV_m = V_{m+1}T_m \quad (7)$$

where T_m is a bidiagonal matrix with the scalars σ_j and λ_j . Therefore, V_m should be orthogonalized using a QR factorization :

$$V_{m+1} = Q_{m+1}R_{m+1} \quad (8)$$

As we show in section 2.2.1 and following the distribution of vectors in Figure 3.(b), the v_i s are distributed in block of consecutive rows between all the processors; However, their overlapped regions are not duplicated between neighboring processors. Thus, to perform the QR factorization, we use an algorithm introduced by Sameh [35] with a parallel implementation provided by Sidje [37]. Recently, a new approach called TSQR has been proposed by Demmel et al. [17] which aims primarily to minimize the communications and better use the BLAS kernel operations. Their current algorithm used in [27] is implemented with POSIX threads and could only be used when a whole matrix V_m is available in the memory of one SMP node. Nevertheless, the algorithm is SPMD-style and should be easily ported to distributed memory nodes.

So far, at the end of the factorization, we get an orthogonal basis Q_{m+1} implicitly represented as a set of orthogonal reflectors. The matrix R_{m+1} is available in the memory of the last processor. To perform the minimization step in GMRES, we produce an Arnoldi-like relation [32, Section 6.3] using equation (7) and (8)

$$M^{-1}AV_m = Q_{m+1}R_{m+1}T_m = Q_{m+1}\bar{G}_m \quad (9)$$

Hence the matrix \bar{G}_m is in Hessenberg form and the new approximate solution is given by $x_m = x_0 + V_m y_m$ where the vector y minimizes the function J defined by

$$J(y) = \|\beta e_1 - \bar{G}_m y_m\|, \quad \beta = \|r_0\| \quad e_1 = [1, 0, \dots, 0]^T \quad (10)$$

The matrix \bar{G}_m is in the memory of the last processor. Since $m \ll n$, this least-square problem is sequentially and easily solved using a QR factorization of G_m . The outline of this GMRES algorithm with the Newton basis can be found in [18].

2.2.1 Parallel processing of the preconditioned Newton basis

In this section, we generate the krylov vectors v_i , ($i = 0, \dots, m$) of V_{m+1} from the equation (5). Consider the partitioning of the Figure 1 with a total of p subdomains. At this point, we assume that the number of processors is equal to the number of subdomains. Thus, each processor computes the sequence of vectors $v_k^{i+1} = (M^{-1}A - \lambda_{i+1}I)v_k^i$ where v_k^i is the set of block rows from the vector v^i owned by the processor P_k . The kernel computation reduces to two major operations $z = Ax$ and $y = M^{-1}z - \lambda x$. For these operations, we consider the matrices and vector distribution on Figures 2 and 3. On each process P_k , the overlapped submatrix is zeroed. We then have the matrix B_i in Figure 2. In Figure 3.(b), the distribution for the vector x is plotted. The overlapped parts are repeated on all processes. Hence, for each subvector x_k on process P_k , x_k^u and x_k^d denote respectively the overlapped parts with x_{k-1} and x_{k+1} . The pseudocode for the matrix-vector product $y = Ax$ follows then in Algorithm 1.

Now we consider the matrix distribution in Figure 3.(a) for the second operation $y = M^{-1}z$. According to the relation (3), each process k solves locally the linear system $A_k t_k = z_k$ for t_k followed by a product $C_k t_k^d = y_k^d$ with the overlapped matrix C_k . However, the current process k should receive first the overlapped part of y_{k-1}^d from the process $k-1$. We thus have the Algorithm 2 for the application of the preconditioner M^{-1} to a vector z . The form of the M^{-1} operator produces a data dependency between neighboring processors. Hence the computation of a single vector $v_j = M^{-1}Av_{j-1}$ is sequential overall the processors. However since the v_j are computed one after another, a process can start to compute its own part of the vector v_j even if the whole previous vector v_{j-1} is not available. We refer to this as a *pipeline parallelism* since the vectors v_j are computed across all the processors as in a pipeline. This would not be possible with the Arnoldi process

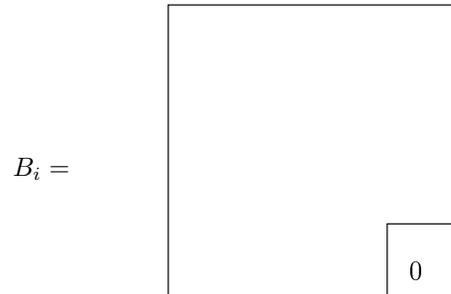


Figure 2: Distribution for the matrix-vector product

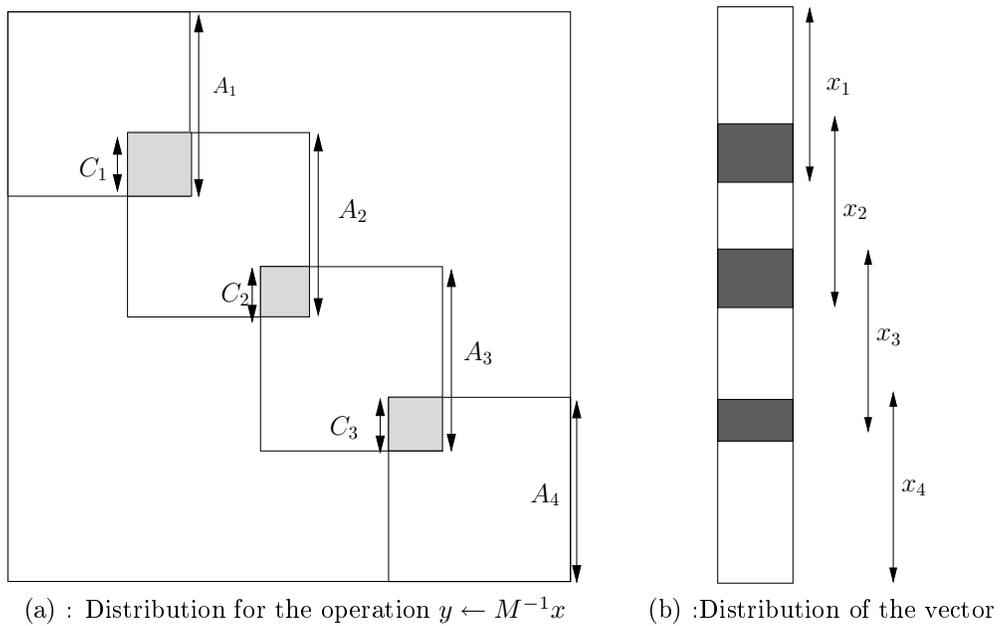


Figure 3: Matrix and vector splittings

Algorithm 1 $z = Ax$

```

1: /* Process  $P_k$  holds  $B_k, x_k$  */
2:  $k \leftarrow \text{myrank}()$ 
3:  $z_k \leftarrow B_k x_k$ ; /* local matrix-vector product */
4: if  $k < p$  then
5:   Send  $z_k^d$  to process  $P_{k+1}$ 
6: end if
7: if  $k > 1$  then
8:   Receive  $z_{k-1}^d$  from process  $P_{k-1}$ 
9:    $z_k^u \leftarrow z_k^u + z_{k-1}^d$ 
10: end if
11: /* Communication for the consistency of overlapped regions */
12: if  $k > 1$  then
13:   Send  $z_k^u$  to process  $P_{k-1}$ 
14:   Receive  $z_{k+1}^u$  from process  $P_{k+1}$ 
15:    $z_k^d \leftarrow z_k^u + z_{k+1}^u$ 
16: end if
17: return  $z_k$ 

```

as all the global communications introduce synchronizations points between the processors and prevent the use of the pipeline; see for instance the data dependencies implied by this process in Erhel [18]. To better understand the actual *pipeline parallelism*, we recall here all

Algorithm 2 $y = M^{-1}z$

```

1: /* Process  $P_k$  holds  $A_k, z_k$  */
2:  $k \leftarrow \text{myrank}()$ 
3: if  $k > 1$  then
4:   Receive  $y_{k-1}^d$  from process  $P_{k-1}$ 
5:    $z_k^u \leftarrow y_{k-1}^d$ 
6: end if
7: Solve local system  $A_k y_k = z_k$  for  $y_k$ 
8: if  $k < p$  then
9:    $y_k^d = C_k y_k^d$ 
10:  Send  $y_k^d$  to process  $P_{k+1}$ 
11: end if
12: /* Communication for the consistency of overlapped regions */
13: if  $k > 1$  then
14:   Send  $y_k^u$  to process  $P_{k-1}$ 
15: end if
16: if  $k < p$  then
17:   Receive  $y_{k+1}^u$  from process  $P_{k+1}$ 
18:    $y_k^d = y_{k+1}^u$ 
19: end if
20: return  $y_k$ 

```

the dependencies presented in [39]. For the parallel matrix-vector product in Algorithm 1, if we set $z = h(x) = Ax$, then $z_k = h_k(x_{k-1}, x_k, x_{k+1})$ and the Figure 4.(a) illustrates these dependencies. For the preconditioner application $y \leftarrow M^{-1}z$ as written in Algorithm 2, we

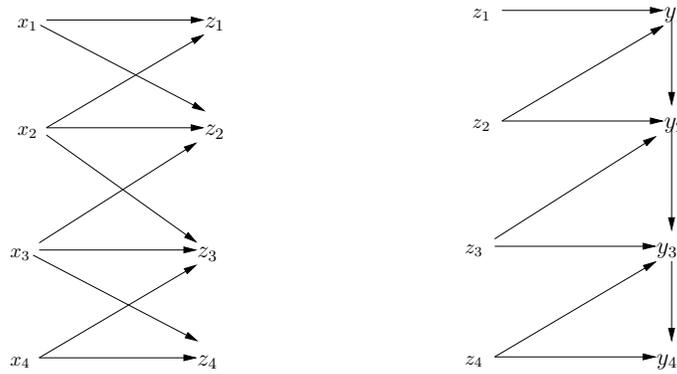
set $y = g(z) = M^{-1}z$, then $y_k = g_k(y_{k-1}, z_k, z_{k+1})$ and we have the dependencies depicted on Figure 4.(b). Finally, for the operation $y \leftarrow M^{-1}Ax$, if $y = f(x) = AM^{-1}x = h \circ g(x)$, then $y_k = f_k(y_{k-1}, x_k, x_{k+1}, x_{k+2})$ and we combine the two graphs to have the dependencies on Figure 5.a. Hence the dependency x_{k-1} is annihilated by that of y_{k-1} . In the pipeline flow computation $v^{i+1} = f(v^i)$, the dependency x_{k+2} in Figure 5.b delays the computation of the v_k^{i+1} until the subvector v_{k+2}^i is available. If there is a good load balancing between all the processes and if τ denotes a time to compute a subvector x_k , then the time to compute one vector is

$$t(1) = p\tau \tag{11}$$

and the time to compute m vectors of the basis is

$$t(m) = p\tau + 3(m - 1)\tau \tag{12}$$

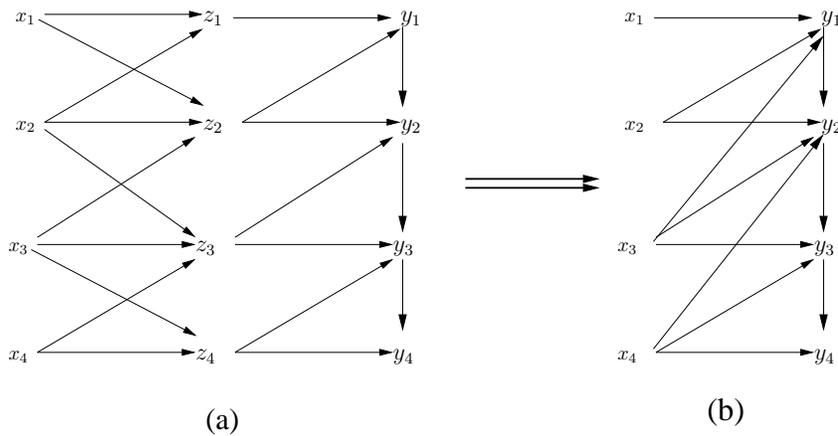
Finally, the first vector is available after $p\tau$ and then, a new vector is produced every 3τ (see Figure 6). The efficiency of the algorithm is therefore equal to $\frac{1}{3}$.



(a): Graph for the matrix-vector

(b): Graph when applying M^{-1}

Figure 4: Dependency graphs



(a)

(b)

Figure 5: Dependency graph for $y = M^{-1}Ax$

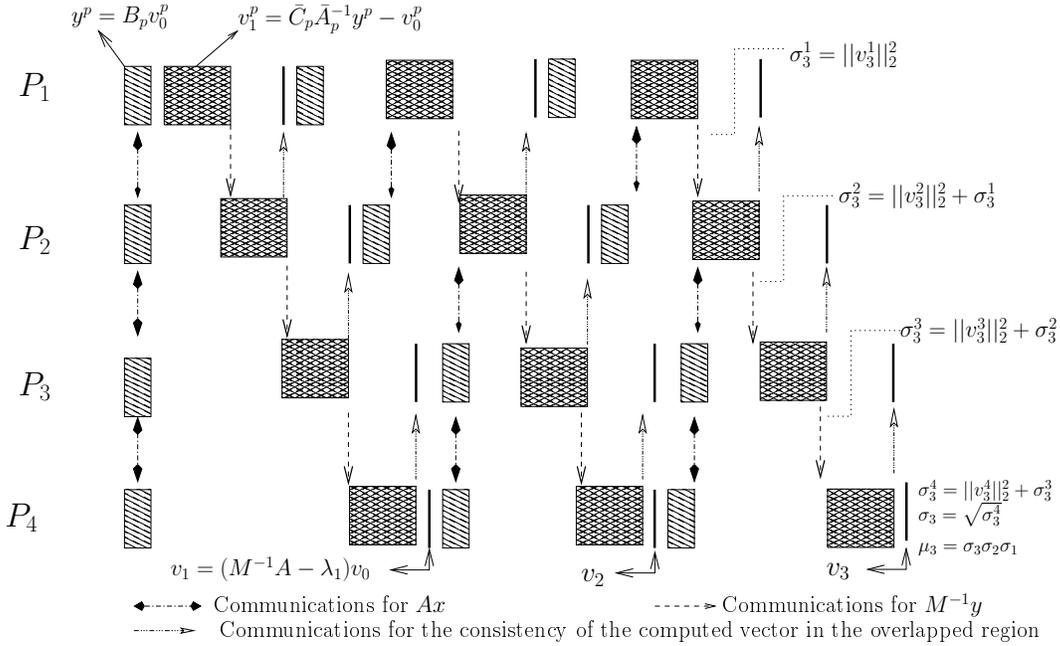


Figure 6: Double recursion during the computation of the Krylov basis

2.2.2 Computation of shifts and scaling factors

So far, we have not yet explained how to compute the shifts λ_i and the scaling factors σ_i of the equation (5) and (6).

As suggested by Reichel [31], the values λ_i should be the m eigenvalues of greatest magnitude of the matrix $M^{-1}A$ in order to have a well-conditioned Krylov basis. Since these eigenvalues can not be cheaply computed, we perform one iteration of classical GMRES(m) with the Arnoldi process. Hence, the m eigenvalues computed from the output Hessenberg matrix are the approximate eigenvalues of the preconditioned global matrix [33, Section 3]. These eigenvalues are arranged using the Leja ordering while grouping together two complex conjugate eigenvalues. Recently, Philippe and Reichel [29] proved that, in some cases, roots of the Chebychev polynomials can be used efficiently.

The scalars σ_i are determined without using global reduction operations. Indeed, such operations introduce synchronization points between all the processes and consequently destroy the pipeline parallelism. The Algorithms 1 and 2 compute in some extent the sequence $v^i = (M^{-1}A - \lambda_i I)v^{i-1}$. We are interested in the scalars $\alpha_i = \|v^i\|$. For this purpose, on process k , we define by \hat{v}_k^i the subvector v_k^i without the overlapped part v_k^{i-u} . In the Algorithm 2, after the line 8, an instruction should be added to compute the local sum $\alpha_i^k = \|\hat{v}_k^i\|^2$. The result is sent to the process $k+1$ at the same time as the overlapped subvector at line 10. The next process $k+1$ receives the result and adds it to its own contribution. This is repeated until the last process p . Then at this step, the scalar $\alpha_i^p = \sqrt{\alpha_i^p}$ is the norm $\|v^i\|$. An illustration is given in the Figure 6 during the computation of the vector v_3 .

3 Parallel computation of local systems in subdomains

During each application of the operator M^{-1} , several linear systems $A_k y_k = z_k$ should be solved for y_k with different values of z_k (see Algorithm 2 line 7). The matrices $A_k, k = (1 \dots p)$ are sparse and nonsymmetric as the global matrix A . Hence, during the whole iterative process and inside each subdomain, a sparse linear system with different right hand sides is being solved. On present supercomputers with SMP (Symmetric Multiprocessing) nodes and with the availability of numerous parallel third-party solvers, it becomes also natural to introduce another level of parallelism within each subdomain. So, in this section, we give first the motivations of using this second level of parallelism. Then we discuss about the appropriate solver to use in subdomains. Finally we give the main steps of our parallel solver with these two levels of parallelism.

3.1 Motivations for two levels of parallelism

In domain decomposition methods, the classical approach is to assign one or several subdomains to a unique processor. In terms of numerical convergence and parallel performance, this approach has some limitations, as pointed out in [22, 24]:

- For some difficult problems, such as those arising from non-elliptic operators, the number of iterations tends to increase with the number of subdomains. Moreover, with the form of the multiplicative Schwarz method, each correction of the residual vector should go through all the subdomain. It is therefore essential to keep this number of subdomains small in order to provide a robust iterative method. By doing this, the subsystems are getting large and could not be solved efficiently with only one processor in a subdomain. In this situation, the second level of parallelism is improving the numerical convergence of the method while providing fast and effective solutions in subdomains.
- For modern computers, each node is usually composed of several processors (see Figure 7.a) accessing more or less the same memory space. For one level of parallelism, we assume that the memory allocated to one processor is not sufficient for one subdomain. Then it becomes necessary to schedule for more nodes in order to access more memory. In the case of one level of parallelism, only one processor will be working (The processor P_0 for instance in Figure 7.b). With the new level of parallelism, all the processors in the node are solving the subsystem in parallel (Figure 7.c). Moreover, the subsystem could be further distributed on more than one node if the whole memory is not sufficient (Figure 7.d).

3.2 Local solver in subdomains : Accuracy and parallelism issues

The sparse linear systems associated to subdomains are in the scope of the preconditioner operator M^{-1} . Therefore, the solutions of these systems can be obtained more or less accurately:

- An accurate solution, with a direct solver for instance, leads to a powerful preconditioner, thus accelerating the convergence of the global iterative method. This approach is usually known as hybrid direct/iterative technique. The drawback here is the significant memory needed in each subdomain. Indeed, not only the submatrix A_k should be stored but also the L_k and U_k matrices from the LU factorization. With the presence of fill-in during numerical factorization, the size of those factors can increase significantly compared to the size of the local matrix.

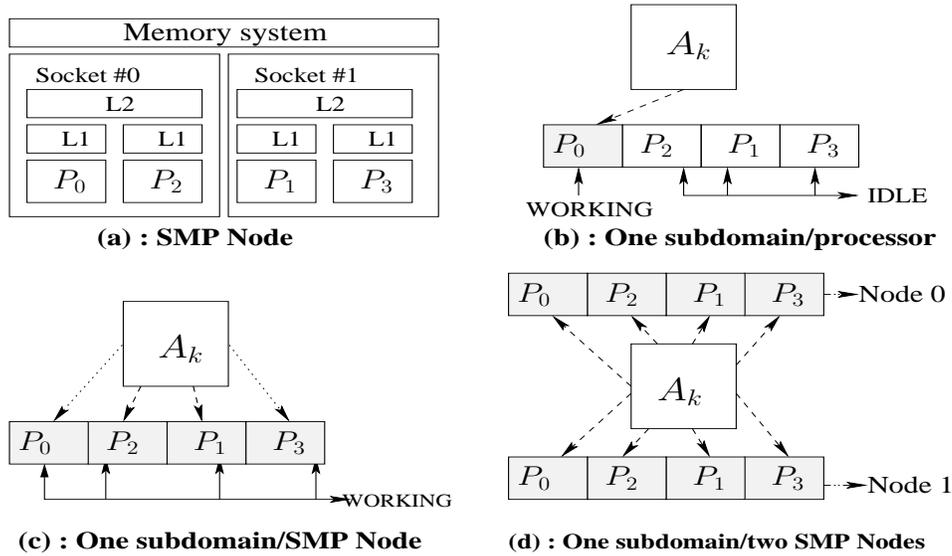


Figure 7: Distribution of subdomains on SMP nodes

- The problem of memory can be overcome by using an incomplete (*ILU*) factorization of local matrices, thus producing fast solutions in subdomains with fair accuracy. In this case, the global iterative method is more likely to stagnate for some complex problems. On another side, an iterative method could be used inside each subdomain. However, a large difference in the accuracy of local solutions could change the preconditioner from one iteration to another leading again to a stagnation.

So far, only the numerical accuracy and the memory issues have been discussed but not the underlying parallel programming model. Beforehand, we assume that a parallel direct solver is used within each computing node when there is enough memory, otherwise an incomplete solver is applied. On SMP nodes with processors sharing the same memory, solvers that distribute tasks among computing units by using thread model are of natural use (Spooles [3], PaStiX [25], PARDISO [36]). On the other side, there are high performance sparse solvers that use the message-passing model. Hence, they are intended primarily for distributed memory processors. However, on SMP nodes, this approach can be very efficient with the availability of intranode communication channels for the underlying message passing interface (MPI) [12, 21]. These channels are appropriate for the large amount of communications introduced between processors during the solution in subdomains; see for instance the comparative study in [2, Sect. 5] depicts a volume and size of these communications. As these messages are of small size, a good data transfer can be obtained inside SMP nodes [11, 30]. For these reasons and the fact that we have the flexibility to split a subdomain on more than one SMP node, we use primarily distributed direct solvers in subdomains for the second level of parallelism.

3.3 Main steps of the parallel solver

The solver goes through almost the same steps whether a direct solver (*LU* factorization) or an incomplete *ILU* factorization is used for the subsystems:

1. **Initialization** : During the first step, the global matrix A is permuted in block diagonal form by the host processor (see Figure 1). After that, the local matrices A_k and C_k should be distributed to other processors. If a distributed solver is used in subdomains, then the processors are dispatched in multiple communicators. For instance, the Figure 8 shows a distribution of four SMP nodes with four processors around two levels of MPI communicators. The first level is intended for the communication across the subdomains. Hence in this communicator, the submatrices are distributed to the *level 0* processors (i.e P_{k0} where $k = 0 \dots d - 1$ and d the number of subdomains). Then for each subdomain k , a second communicator is created between the *level 1* processors to manage communications inside the subdomains (i.e P_{kj} where $j = 0 \dots p_k - 1$ and p_k the number of processors in the node k).

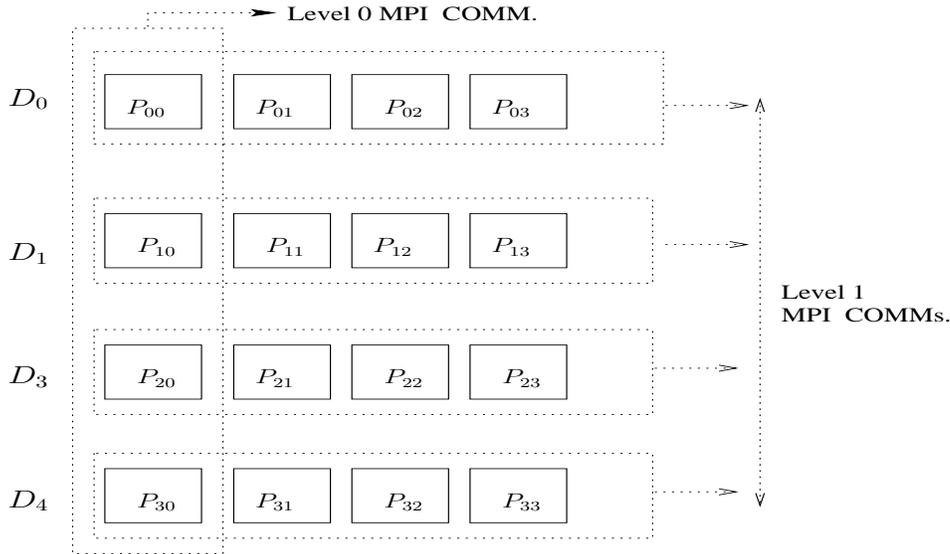


Figure 8: MPI communicators for the two levels of parallelism

2. **Setup** : In this phase, the symbolic and numerical factorization are performed on submatrices A_k by the underlying local solver. This step is purely parallel across all the subdomains. At the end of this phase, the factors L_k and U_k reside in the processors responsible for the subdomain k . This is totally managed by the underlying local solver. Prior to this phase, a preprocessing step can be performed to scale the elements of the matrix.
3. **Solve** : This is the iterative phase of the solver. With an initial guess x_0 and the shifts λ_i , the solver computes all the equations (5-10) as outlined here:
 - (a) Pipeline computation of V_{m+1}
 - (b) Parallel QR factorization $V_{m+1} = Q_{m+1}R_{m+1}$
 - (c) Computation of \bar{G}_m such that $M^{-1}AV_m = Q_{m+1}\bar{G}_m$
 - (d) Solution of least-square problem (10) for y_m ;
 - (e) Computation of $x_m = x_0 + V_my_m$

The convergence is reached when $\|b - Ax\| < \epsilon\|b\|$ otherwise the process restart with $x_0 = x_m$. When two levels of parallelism are used during the computation of V_m , *level 1* processors perform multiple backward and forward sweeps in parallel to solve local systems. After the parallel QR factorization in step 3b, the explicit Q factor is never formed explicitly. Indeed, only the unfactored basis V is used to apply the new correction as shown in step 3e. Nevertheless, a routine to form this factor is provided in the solver for general purposes.

4 Numerical experiments

In this section, we perform several experiments to demonstrate the parallel performance and the numerical robustness of our parallel solver on a wide range of real problems. In section 4.1, we give first the software and the hardware architecture on which the tests are done. The sections 4.3 and 4.4 provide intensive results in terms of scalability and robustness on several tests cases. Those cases listed in Table 1 are taken from various sources, ranging from public repositories to industrial softwares

4.1 Software and hardware framework

The solver is named GPREMS¹ (GMRES PRECONDITIONED BY MULTIPLICATIVE SCHWARZ). It is intended to be deployed on distributed memory computers that communicate through message passing (MPI). The parallelism in subdomains is based either on message passing or threads model depending on the underlying solver. The whole library is built on top of PETSc (Portable, Extensible Toolkit for Scientific Computation) [8, 9]. The motivation of this choice is the uniform access to a wide range of external packages for the linear systems in subdomain. Moreover, the PETSc package provide several optimized functions and data structures to manipulate the distributed matrices and vectors. A Figure 9 gives the position

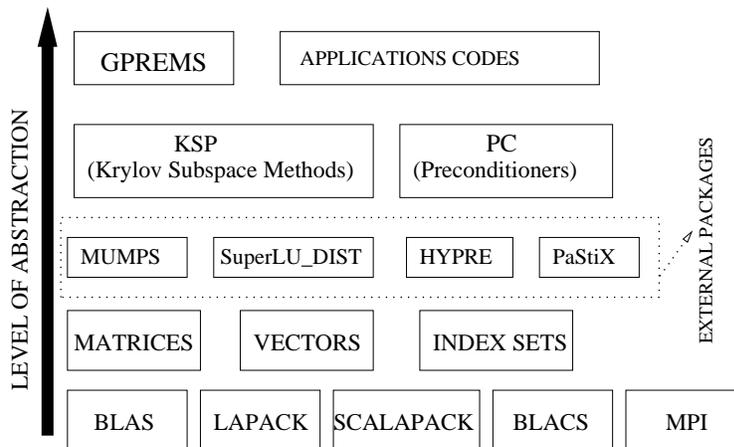


Figure 9: GPREMS library in PETSc

of GPREMS in the abstraction layer of PETSc. We give only the components that are used by GPREMS. For a complete reference on PETSc environment, please refer to [8]. Although

¹A public license will be released soon

GPREMS uses the PETSc environment, it is not distributed as a part of PETSc libraries. Nevertheless, the library is configured easily once a usable version of PETSc is available on the targeted architecture.

All the tests in this paper are performed on the IBM p575 SMP nodes connected through the Infiniband DDR network. Each node is composed of 32 Power6 processors sharing the same L3 cache. A Power6 processor is a dual-core 2-way SMT with a peak frequency at 4.7 GHz. Figure 10 (obtained with *hwloc*[10]) shows the topology of the first eight processors in one compute node. A total of 128 nodes is available in this supercomputer (named *Vargas*²) which is part of the French CNRS-IDRIS computing facilities.

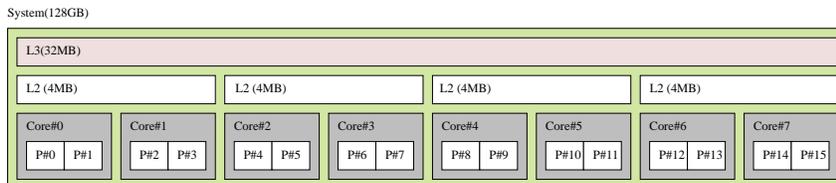


Figure 10: Vargas Architecture

4.2 Test matrices

This section gives the main characteristics of the set of problems under study. The first set is taken from the University of Florida (UFL) Sparse Matrix Collection [15]. In this set, we choose the problems MEMCHIP and PARA-4 respectively from Freescale and Schenk_ISEI directories; main characteristics of which are given in Table 1. The MEMCHIP problem

Table 1: General properties of the two CFD cases

Matrix Name	MEMCHIP	PARA_4	CASE_004	CASE_017
Order	2,707,524	153,226	7,980	381,689
No. entries (nnz)	14,810,202	2,930,882	430,909	37,464,962
Entries per row	5	19	53	98
Pattern symmetry	91%	100%	88%	93%
2D/3D problem	No	yes	2D	3D
Source	UFL (K. Gullapalli)	UFL (O. Schenk)	FLUOREM	FLUOREM

arise from circuit simulation field. The density of the associated matrix is rather small. The number of entries is about $1.4 \cdot 10^6$ but the nonzeros elements per row reveals a rather small density. Hence, this problem is easily solved with GPREMS using an incomplete factorization in subdomains. We present the results with this case in section 4.3. The next matrix is PARA-4 coming from 2D semiconductor device simulations. Compared to the previous problem, the geometry of this problem makes it difficult to solve. Hence, we choose this matrix to show the robustness of our solver when it is combined with a direct solver in subdomains. Again, the results on this are shown in section 4.3.

The second set contains two test matrices (CASE_004 and CASE_017) provided by FLUOREM [20], a fluid dynamics software publisher. In those cases, the sparse matrix corresponds to the global jacobian matrix resulting from the partial first-order derivatives of the discrete steady Euler or Reynolds-averaged Navier-Stokes equations. The derivatives are

²<http://www.idris.fr/su/Scalaire/vargas/hw-vargas.html>

done with respect to the seven conservative fluid variables (mass: ρ , momentum: $[\rho u, \rho v, \rho w]$, energy: ρE , turbulence: $[\rho k, \rho \omega]$ with $k - \omega$); with Q denoting these variables and $\mathcal{F}(Q)$ the divergence of the flux function from the flow equations written in the conservative form, then the steady equation is $\mathcal{F}(Q) = 0$. If Q_h and \mathcal{F}_h are the finite volume discretizations over the computational domain of Q and \mathcal{F} respectively, then the matrix A of the linear system (1) is the global jacobian matrix $J_{\mathcal{F}_h}(Q_h)$. $A \in \mathbb{R}^{n \times n}$ is non-symmetric and indefinite and when the stationary flow is dominated by advection, it corresponds to a strongly non-elliptic operator. The two-level parallelism is then useful in those cases in that, we can increase the number of processors to speed up the overall execution time while providing the guarantee that the number of iterations will remain almost constant. The section 4.4 gives the results with the multilevel approach on the CASE_017 3D case. The matrix CASE_004 is used to show the numerical robustness of GPREMS compared to the additive Schwarz approach implemented in PETSc.

4.3 Convergence and scalability studies on MEMCHIP and PARA-4 cases

We discuss here the results of the numerical convergence of GPREMS on the two problems MEMCHIP and PARA-4. We use respectively $m = 12$ and $m = 24$ as the size of the Newton basis (search directions) for the two cases. The iterative process stops when the relative residual norm $\|b - Ax\|/\|b\|$ is less than 10^{-08} or when 1500 iterations are reached. The total number of iterations is usually a multiple of m as the Newton basis should be generated a priori; hence the convergence is tested only at each external cycle. $ILU(p)$ denotes a use of an incomplete factorization in subdomains with p levels of *fill-in* in factored matrix. The package used to perform this factorization is EUCLID[26] from the HYPRE suite [19]. Some other problems like PARA-4 are more difficult to solve. Hence the use of hybrid direct-iterative solver is necessary to achieve a good convergence in a reasonable CPU time. The MUMPS [1] package is then used at this point. On all cases so far, the relative residual norm of the computed solution is given with respect to the number of iterations. At each time, we compare the convergence history on 4, 8 and 16 subdomains. Following the convergence history, we give the time spent in the various phases setup phase, the iterative loop.

For the MEMCHIP case with about $1.5 \cdot 10^7$ entries, it takes less than 25 cycles to converge as shown on Figure 11. Moreover, the number of iterations tends to decrease with 16 subdomains. The $ILU(0)$ factorization is used in subdomains meaning that no new elements are allowed during the numerical factorization and thus no more space is needed. This has the effect to speed-up the setup phase and the application of the M^{-1} operator in the iterative step. The CPU time of these steps are shown in Table 2 along with the initialization time (Init) and the total time. These results prove the scalability for almost all the steps in GPREMS. Apart from the initialization step which is sequential, all others steps benefit from the first level of parallelism in the method. The setup phase is fully parallel across the subdomains. Here each process responsible for a subdomain performs a sequential factorization on its local matrix. As the size of those submatrices decrease with the number of subdomains, the setup time decreases as well. The iterative part of the solver (given by *Iter. loop time* in Table 2) scales very well in this case. For instance, 300 iterations of GMRES are performed in less than 103 s. when using 4 subdomains in the block-diagonal partitioning. With 8 subdomains, the method requires 312 iterations (26 restarts or cycles) computed in less than 65 s. If the problem is subdivided in 16 subdomains, the method reveals a superlinear speedup, thanks to the rather small number of iterations needed to converge.

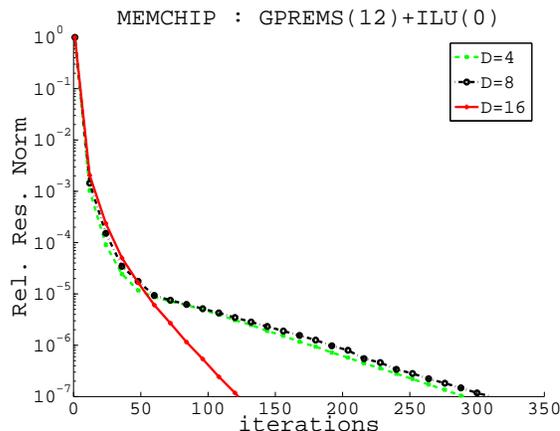


Figure 11: Convergence history with MEMCHIP case on 4, 8 and 16 subdomains

Table 2: MEMCHIP : CPU Time for various number of subdomains

#Processors	#Subdomains	Iterations	CPU Time (s.)					
			Init	Setup	Iter. loop	Time/Iter	Total	
Incomplete solver in subdomains (ILU(0) - EUCLID)								
4	4	300	14.96	3.16	102.31	0.34	120.42	
8	8	312	15.35	1.7	64.35	0.21	81.41	
16	16	132	16.07	0.95	19.98	0.15	37	

Obviously, when the subsystems are solved approximately, the global method loses some robustness. Thus for many 2D/3D cases, the number of iterations gets very large and consequently the overall time needed to converge. We illustrate this behavior on the PARA-4 2D case. An incomplete LU factorization ($ILU(2)$) is first performed on each submatrix to solve the systems induced by the preconditioner operator. The $ILU(2)$ factorization keeps the second-order fill-ins in the factors, leading to a more accurate preconditioner than that with $ILU(0)$. However, this is not good enough to achieve a good convergence rate as in the previous case. For instance in Figure 12.(a), after reaching 10^{-07} , the residual norm tends to stagnate. This can be observed either with 4, 8 or 16 subdomains. On the other hand, it takes a certain time to setup the block matrices for the iterative phase as shown in Table 3. It can be seen that the maximum time overall processors (*Setup*) to get the incomplete factorization of local matrices is twice than the time of the iterative loop. Thus, the method is less than scalable for this case.

Now if a direct solver is used within each subdomain, the method will get a better convergence thanks to accuracy achieved during the application of the preconditioner. This behavior is observed in Figure 12.(b) in which we show the convergence history for the number of subdomains under study. The number of iterations grows slightly when we change from 4 to 8 subdomains. This does not impact the overall time spent in the iterative step as shown in the second part of the Table 3. Indeed, this step reveals a very good speedup as we only need 35 s. with 8 subdomains compared to 45 s. for 4 subdomains. With 16 processors (subdomains), the total time continues to decrease, even with the slight increase of the global number of iterations. It is also worth pointing out the time spent in the setup phase. It takes roughly 5 s. on four processors to perform the symbolic and the numerical

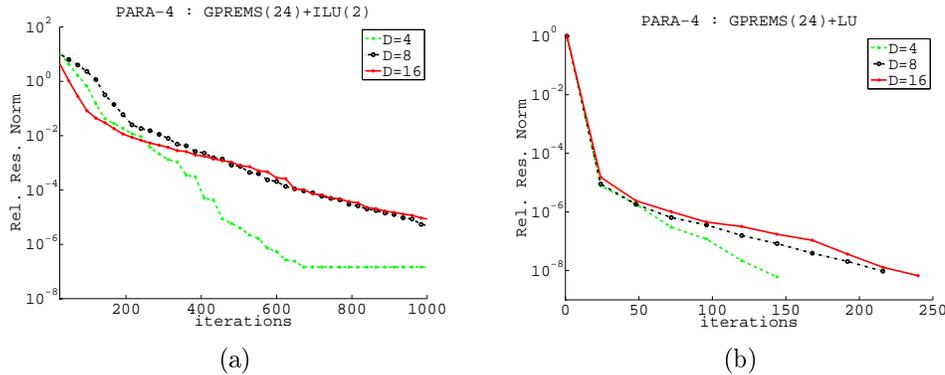


Figure 12: Convergence history with PARA-4 case on 4, 8 and 16 subdomains

factorization of the submatrices. Note that each process holds a local matrix associated to one subdomain. Hence, the setup time is the maximum time overall active processors. When the matrix is splitted into 16 subdomains, the size of these submatrices is rather small and the setup time drops to 1 s. This is almost 400 times smaller than that with the ILU(2) factorization. The main reason is the algorithms used by the two methods. In MUMPS, almost all the operations are done on block matrices (the so-called frontal matrices) *via* calls to level 2 and level 3 BLAS functions. Therefore, the different levels of the cache memory are used efficiently during the numerical factorization which is the most-time consuming step. Moreover, a multithreaded BLAS is linked to MUMPS to further speedup the method on SMP nodes. For this reason and the rather small size of submatrices, the second level of parallelism is not applied up to this point.

Table 3: **PARA-4** : CPU Time for various number of subdomains and various types of solvers in subdomains

#Processors	#Subdomains	Iterations	CPU Time (s.)				
			Init	Setup	Iter. loop	Time/Iter	Total
Incomplete solver in subdomains (ILU(2) - EUCLID)							
4	4	1008	1.33	418.17	654.45	0.65	1073.95
8	8	1008	1.31	409.68	599.64	0.59	1010.63
16	16	1008	1.51	398.61	542.75	0.54	942.87
Direct solver in subdomains (LU - MUMPS)							
4	4	144	1.32	5.12	45.43	0.32	51.87
8	8	216	1.26	1.67	35.06	0.16	37.98
16	16	240	1.37	0.73	29.21	0.12	31.31

4.4 Benefits of the two-levels parallelism

Now, let us consider the FLUOREM problem CASE_017 listed in Table 1. The geometry of this 3D case is a jet engine compressor. Although the turbulent variables are frozen during the differentiation, this test case is difficult to solve as shown in a short comparative study involving some distributed linear solvers [28]. From this study, solvers that combine hybrid direct-iterative approaches like GPREMS provide an efficient way to deal with such problems. However, the number of subdomains should be kept small to provide a good

convergence. Therefore the two-levels of parallelism introduced in section 3.2 provide the way to do this efficiently even with a large number of processors. In Table 4, we report the benefit of using this approach on the above-named 3D case. We keep 4 and 8 subdomains and we increase the total number of processors. As a result, almost all the steps in GPREMS get a noticeable speedup. Typically, with 4 subdomains, when only one processor is active, the time to setup the block matrices is almost 203 s. and the time spent in the iterative loop is 622 s. Moving to 8 active processors decreases these times to 59 s. and 181 s. respectively. The same observation can be done when using 8 subdomains, in which case the overall time drops from 540 s. to 238 s. Note that the overall time includes the first sequential step that permute the matrix in block-diagonal form and distribute the block matrices to all active processors. From 32 to 64 processors, the overall time decreases from 311 s. to 237 s. indicating a little speedup. This is mainly due to the total number of iterations. Indeed, if we look at the time spent in each iteration, we observe a very good efficiency when moving from 32 to 64 active processors.

Table 4: Benefits of the two-levels of parallelism for various phases of GPREMS on CASE_17

#Subdomains	#Processors	Iterations	CPU Time (s.)				
			Init	Setup	Iter. loop	Time/Iter	Total
Direct solver in subdomains (MUMPS)							
4	4	128	70.78	203.67	622.94	4.87	897.39
	8	128	70.77	125.77	411.42	3.21	607.96
	16	128	69.79	73.15	280.41	2.19	423.35
	32	128	70.77	59.44	181.45	1.42	311.66
8	8	256	69.91	71.73	399.34	1.56	540.98
	16	256	70.36	42.99	343.25	1.34	456.59
	32	256	71.61	28.72	206.7	0.81	307.02
	64	256	71.85	20.85	144.79	0.57	237.49

Another motivation for the two levels of parallelism is to better use the allocated processors. Indeed, when using only one processor per subdomain, the factored block matrices could not fit the local memory of the processor. Hence, it becomes necessary to schedule for more processors in the SMP node in order to access their memory. This approach leaves several processors idle during the whole computation. For instance, with 4 subdomains, the largest local matrix overall processors has 108,150 rows/columns and contains almost $12 \cdot 10^6$ entries. To perform the numerical factorization, MUMPS requires at least 11 GB memory on each of the four processes. When a processor is scheduled on our target cluster, only 3.5 GB is available. Therefore, 3 more processors should be added to each existing process to have enough memory for the setup phase and all others memory allocations. It means that a total of 12 processors are idle during the whole computation. With the second level of parallelism, all these processes work, each on a piece of a particular subdomain.

4.5 Numerical robustness of GPREMS

In this section, our aim is to further show the robustness of the GPREMS solver on CASE_004 listed in Table 1. The matrix of this 2D test case corresponds to the global jacobian of the discrete steady Euler equation [20]. The two turbulent variables are kept during the differentiation predicting a difficult convergence for any iterative method. The reciprocal condition number is estimated to be $3.27 \cdot 10^{-06}$. The matrix is permuted in 4 block diagonal matrices such that each local matrix has approximately 2,000 rows/columns

and 120,000 nonzero entries. For these relatively small block matrices, we choose an incomplete local solver in subdomains when applying the M^{-1} operator. We also allow 5 degrees of fill-ins ($ILU(5)$) during the incomplete numerical factorization. The Figure 13 gives the history of the preconditioned residual norm at each restart for 16 and 32 krylov vectors. This norm drops to 10^{-11} in less than 200 iterations for GPREMS(16) with $ILU(5)$ in subdomains. When the size of the Krylov subspace vary to 32, the residual vector reaches 10^{-9} in less than 100 iterations. This is almost the same behavior when a direct solver is used within the subdomains in which case the overall number of iterations is 112 and 96 respectively for 16 and 32 Krylov vectors. Then we compare with the restricted additive Schwarz

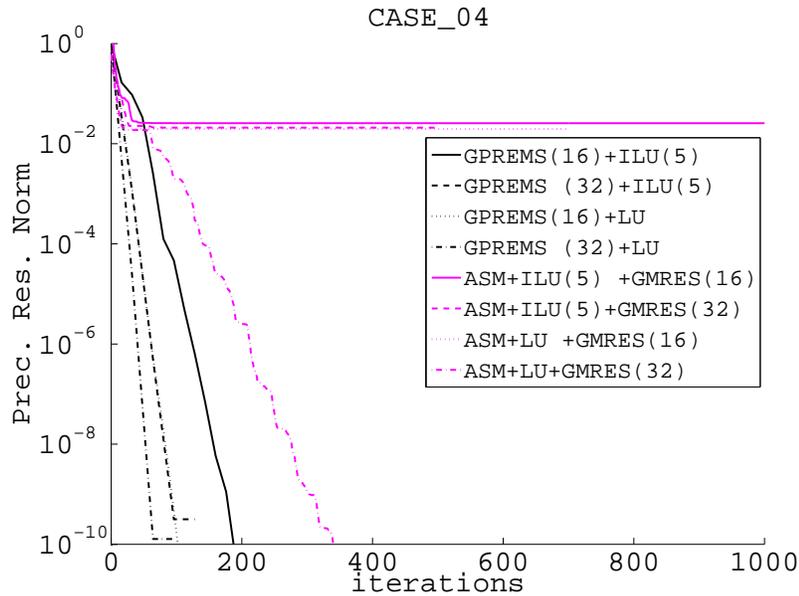


Figure 13: The multiplicative Schwarz approach (GPREMS) compared to the restricted additive Schwarz (ASM) on CASE_004; 16 and 32 vectors in the Krylov subspace at each restart; 4 subdomains (block diagonal partitioning in GPREMS and Parnetis for ASM); Local systems solved with different accuracies ($ILU(5)$ and LU)

method used as a preconditioner for GMRES. The tests are done with the implementation provided in PETSc release 3.0.0-p2. Here, the Krylov basis is built with the modified Gram-Schmidt (MGS) method and the input matrix is partitioned in 4 subdomains using ParMETIS. In Figure 13, we depict again the preconditioned residual norm with respect to the number of iterations. The curves are given in magenta with the same line specifications as for GPREMS. The main observation is that the iterative method stagnates from the first restart when using $ILU(5)$ in subdomains (the magenta plain and dash curves). With a direct solver, it is necessary to form 32 Krylov vectors at each restart in order to achieve a good convergence (as shown by the magenta dash-dotted curve).

5 Concluding remarks

In this paper, we give an implementation of a parallel solver for the solution of linear systems on distributed-memory computers. This implementation is based on an hybrid technique

that combines a multiplicative Schwarz preconditioner with a GMRES accelerator. A very good efficiency is obtained using a Newton basis GMRES version; hence we define two levels of parallelism during the computation of the orthonormal basis needed by GMRES : The first level is expressed through pipeline operations across all the subdomains. The second level uses a parallelism inside third-party solvers to build the solution of subsystems induced by the domain decomposition. The numerical experiments demonstrate the efficiency of this approach on a wide range of test cases. Indeed, with the first-order parallelism, we show a good efficiency when using either a direct solver or an incomplete LU factorization. For the second level used within the subdomains, the gain obtained is two-fold: the convergence is guaranteed when the number of processors grows as the number of subdomains remains the same. The global efficiency increases as we add more processors in subdomains. Moreover, on some difficult problems, the use of a direct solver in subdomains implies to access the whole memory of a SMP node. Therefore, this approach keeps busy all the processors of the node and enables a good usage of allocated computing resources. Finally, we show that this solver can be competitive on some test cases with respect to other overlapping domain decomposition preconditioners.

However, more work needs to be done to achieve very good scalability on massively parallel computers. Presently, an attempt to increase the number of subdomains up to 32 increases the number of iterations as well. The main reason is the form of the multiplicative Schwarz method as each correction of the residual vector should go through all the subdomains. A first attempt is to use more than two levels of splitting but the efficiency of such approach is not always guaranteed, specifically if the linear system arises from non-elliptic partial differential equations. An ongoing work is to keep the two-levels of splitting for the preconditioner operator and then to further accelerate the GMRES method with spectral information gathered during the iterative process.

References

- [1] AMESTOY, P. R., DUFF, I. S., L'EXCELLENT, J.-Y., AND KOSTER, J. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM Journal on Matrix Analysis and Applications* 23, 1 (2001), 15–41.
- [2] AMESTOY, P. R., DUFF, I. S., L'EXCELLENT, J.-Y., AND LI, X. S. Analysis and comparison of two general sparse solvers for distributed memory computers. *ACM Transactions on Mathematical Software* 27 (2000), 2001.
- [3] ASHCRAFT, C., AND GRIMES, R. Spooles: An object-oriented sparse matrix library. In *Proceedings of the 9th SIAM Conference on Parallel Processing for Scientific Computing* (1999).
- [4] ATENEKENG-KAHOUE, G.-A. *Parallélisation de 'GMRES' préconditionné par une itération de Schwarz multiplicatif*. PhD thesis, University of Rennes 1 and University of Yaounde 1, 2008.
- [5] ATENEKENG-KAHOUE, G.-A., GRIGORI, L., AND SOSONKINA, M. A partitioning algorithm for block-diagonal matrices with overlap. *Parallel Computing* 34, 6-8 (2008), 332–344.
- [6] ATENEKENG-KAHOUE, G.-A., KAMGNIA, E., AND PHILIPPE, B. An explicit formulation of the multiplicative schwarz preconditioner. *Applied Numerical Mathematics* 57, 11-12 (2007), 1197 – 1213.

-
- [7] BAI, Z., HU, D., AND REICHEL, L. A Newton basis GMRES implementation. *IMA J Numer Anal* 14, 4 (1994), 563–581.
- [8] BALAY, S., BUSCHELMAN, K., EIJKHOUT, V., GROPP, W. D., KAUSHIK, D., KNEPLEY, M. G., MCINNES, L. C., SMITH, B. F., AND ZHANG, H. PETSc users manual. Tech. Rep. ANL-95/11 - Revision 3.0.0, Argonne National Laboratory, 2008.
- [9] BALAY, S., BUSCHELMAN, K., GROPP, W. D., KAUSHIK, D., KNEPLEY, M. G., MCINNES, L. C., SMITH, B. F., AND ZHANG, H. PETSc Web page, 2009. <http://www.mcs.anl.gov/petsc>.
- [10] BROQUEDIS, F., CLET-ORTEGA, J., MOREAUD, S., FURMENTO, N., GOGLIN, B., MERCIER, G., THIBAUT, S., AND NAMYST, R. hwloc: a generic framework for managing hardware affinities in HPC applications. In *Proceedings of the 18th Euromicro International Conference on Parallel, Distributed and Network-Based Processing (PDP2010), Pisa, Italia* (2010).
- [11] BUNTINAS, D., AND MERCIER, G. Data transfers between processes in an SMP system: performance study and application to MPI. In *Proceedings of the 2006 International Conference on Parallel Processing (ICPP 2006), IEEE Computer Society* (2006), pp. 487–496.
- [12] BUNTINAS, D., MERCIER, G., AND GROPP, W. Implementation and evaluation of shared-memory communication and synchronization operations in mpich2 using the nemesis communication subsystem. *Parallel Computing* 33, 9 (2007), 634–644. Selected Papers from EuroPVM/MPI 2006.
- [13] CAI, X.-C., AND SAAD, Y. Overlapping domain decomposition algorithms for general sparse matrices. *Numerical Linear Algebra with Applications* 3, 3 (1996), 221–237.
- [14] CARVALHO, L., GIRAUD, L., AND MEURANT, G. Local preconditioners for two-level non-overlapping domain decomposition methods. *Numerical Linear Algebra with Applications* 8, 4 (2001), 207–227.
- [15] DAVIS, T. The university of Florida sparse matrix collection, 2010. Submitted to ACM Transactions on Mathematical Software.
- [16] DE STURLER, E. A parallel variant of GMRES(m). In *Proceedings of the 13th IMACS World Congress on Computation and Applied Mathematics, Dublin, Ireland* (1991), J. J. H. Miller and R. Vichnevetsky, Eds., Criterion Press.
- [17] DEMMEL, J., GRIGORI, L., HOEMMEN, M., AND LANGOU, J. Communication-optimal parallel and sequential QR and LU factorizations. Tech. Rep. UCB/EECS-2008-89, University of California EECS, May 2008.
- [18] ERHEL, J. A parallel GMRES version for general sparse matrices. *Electronic Transaction on Numerical Analysis* 3 (1995), 160–176.
- [19] FALGOUT, R., AND YANG, U. Hypre: a library of high performance preconditioners. In *Lectures Notes in Computer Science* (2002), vol. 2331, Springer-Verlag, pp. 632–641.
- [20] FLUOREM. The fluorem matrix collection, 2009. LIB0721 2.0 / FP-SA <http://www.fluorem.com>.

-
- [21] GEOFFRAY, P., PRYLLI, L., AND TOURANCHEAU, B. Bip-smp: High performance message passing over a cluster of commodity smps. In *Proceedings of the 1999 ACM/IEEE Conference on Supercomputing* (1999).
- [22] GIRAUD, L., HAIDAR, A., AND PRALET, S. Using multiple levels of parallelism to enhance the performance of domain decomposition solvers. *Parallel Computing In Press* (2010).
- [23] GIRAUD, L., HAIDAR, A., AND WATSON, L. Parallel scalability study of hybrid preconditioners in three dimensions. *Parallel Computing* 34, 6-8 (2008), 363 – 379.
- [24] HAIDAR, A. *On the parallel scalability of hybrid linear solvers for large 3D problems*. PhD thesis, Institut National Polytechnique de Toulouse, 2008.
- [25] HÉNON, P., PELLEGRINI, F., RAMET, P., ROMAN, J., AND SAAD, Y. High Performance Complete and Incomplete Factorizations for Very Large Sparse Systems by using Scotch and PaStiX softwares. In *Eleventh SIAM Conference on Parallel Processing for Scientific Computing* (2004).
- [26] HYSOM, D., AND POTHEN, A. A scalable parallel algorithm for incomplete factor preconditioning. *SIAM Journal on Scientific Computing* 22 (2000), 2194–2215.
- [27] MOHIYUDDIN, M., HOEMMEN, M., DEMMEL, J., AND YELICK, K. Minimizing communication in sparse matrix solvers. In *SC '09: Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis* (New York, NY, USA, 2009), ACM, pp. 1–12.
- [28] NUENTSA WAKAM, D., ERHEL, J., ÉDOUARD CANOT, AND ATENEKENG KAHOU, G.-A. A comparative study of some distributed linear solvers on systems arising from fluid dynamics simulations. In *Parallel Computing: From Multicores and GPU's to Petascale* (2010), vol. 19 of *Advances in Parallel Computing*, IOS Press, pp. 51–58.
- [29] PHILIPPE, B., AND REICHEL, L. On the generation of krylov subspace bases. Tech. Rep. RR-7099, INRIA, 2009.
- [30] RABENSEIFNER, R., HAGER, G., AND JOST, G. Hybrid mpi/openmp parallel programming on clusters of multi-core smp nodes. *Parallel, Distributed, and Network-Based Processing, Euromicro Conference on 0* (2009), 427–436.
- [31] REICHEL, L. Newton interpolation at leja points. *BIT* 30 (1990), 305–331.
- [32] SAAD, Y. *Iterative Methods for Sparse Linear Systems*, second ed. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.
- [33] SAAD, Y., AND SCHULTZ, M. H. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing* 7, 3 (1986), 856–869.
- [34] SAAD, Y., AND SOSONKINA, M. Distributed schur complement techniques for general sparse linear systems. *SIAM Journal on Scientific Computing* 21, 4 (1999), 1337–1356.
- [35] SAMEH, A. Solving the linear leastsquares problem on a linear array of processors. In *High Speed Computer and Algorithm Organization* (1977), D. L. D. Kuck and A. Sameh, Eds., Academic Press, pp. 207–228.

-
- [36] SCHENK, O., AND GÄRTNER, K. Solving unsymmetric sparse systems of linear equations with pardiso. *Journal of Future Generation Computer Systems* 20 (2004), 475–487.
- [37] SIDJE, R. B. Alternatives for parallel krylov subspace basis computation. *Numerical Linear Algebra with Applications* 4, 4 (1997), 305–331.
- [38] SMITH, B., BJØRSTAD, P., AND GROPP, W. *Domain Decomposition, Parallel Multi-level Methods for Elliptic Partial Differential Equations*. Cambridge University Press, 1996.
- [39] SOUOPGUI, I. Parallélisation d'une double recurrence dans 'gmres'. Master's thesis, University of Yaounde 1, 2005.
- [40] WALKER, H. F. Implementation of the 'gmres' method using householder transformations. *SIAM Journal on Scientific and Statistical Computing* 9, 1 (1988), 152–163.



Centre de recherche INRIA Rennes – Bretagne Atlantique
IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399