

# pARMS: a Parallel Version of the Algebraic Recursive Multilevel Solver

Zhongze Li<sup>1,†,§</sup>, Yousef Saad<sup>2,\*;†</sup> and Masha Sosonkina<sup>3,§</sup>

<sup>1</sup> INRIA Projet CAIMAN, 2004 Route des Lucioles, B.P. 93, 06902 Sophia-Antipolis Cedex, France.

<sup>2</sup> University of Minnesota, 200 Union Street S.E., Minneapolis, MN 55455, U.S.A.

<sup>3</sup> University of Minnesota - Duluth, 320 Heller Hall, 10 University Drive, Duluth, MN 55812-2496, U.S.A.

## SUMMARY

A parallel version of the Algebraic Recursive Multilevel Solver (ARMS) is developed for distributed computing environments. The method adopts the general framework of distributed sparse matrices and relies on solving the resulting distributed Schur complement system. Numerical experiments are presented which compare these approaches on regularly and irregularly structured problems. Copyright © 2002 John Wiley & Sons, Ltd.

KEY WORDS: multilevel ILU preconditioner; parallel preconditioning; Schur complement preconditioner; distributed sparse linear systems; domain decomposition; multigrid

## 1. Introduction

Two competing classes of iterative methods can be used for solving large linear systems of equations: multigrid techniques and (preconditioned) Krylov subspace techniques. Methods in the first group are known to be optimal, in terms of overall cost, for certain types of problems. Their implementation requires a hierarchy of (simple) grids and may necessitate specific tailoring for particular applications. Algebraic MultiGrid (AMG) methods have been proposed as an alternative, for solving more general problems, but their overall success, which in theory still relies on an underlying PDE problem, has been somewhat limited. In contrast, preconditioned Krylov methods, using incomplete LU (ILU) preconditioners, are designed to be “general-purpose” methods for solving arbitrary sparse linear systems of equations. In the

---

\*Correspondence to: Y. Saad, Department of Computer Science and Engineering, University of Minnesota, 200 Union Street S.E., Minneapolis, MN 55455, U.S.A.

†E-mail: Zhongze.Li@sophia.inria.fr

‡E-mail: saad@cs.umn.edu

§E-mail: masha@d.umn.edu

Contract/grant sponsor: NSF; contract/grant number: ACI-0000443

Contract/grant sponsor: NSF; contract/grant number: INT-0003274

Contract/grant sponsor: Minnesota Supercomputing Institute

past, users have often sacrificed the speed of multigrid for the better robustness and generality of preconditioned Krylov solvers. However, as the problems become larger the advantage of multilevel approaches can be overwhelming. An attractive middle-ground solution would be a method whose cost scales well with the size of the problem and which is as general purpose as the ILU-Krylov combination.

A number of methods were recently developed to try to fill this gap by introducing concepts borrowed from multilevel methods into incomplete factorizations. ILUM [1] and a few related methods [2, 3] showed that this approach is fairly robust and that it scales well with problem size, in contrast with standard ILU preconditioners. The idea was extended to block versions using essentially domain decomposition type strategies [4, 5, 6]. In [5, 6, 7] it was shown that the block approach is generally more efficient and more robust than a standard ILUT-preconditioned GMRES [8] and than its scalar sibling ILUM. A number of similar techniques have been derived by taking the opposite viewpoint, i.e., by starting from algebraic multigrid solvers, see for example [9, 10]. Other methods have also tried to integrate general purpose approximate inverse preconditioners as smoothers in multilevel techniques [11, 12]. Recently, parallel software packages have emerged which are based on, or include, AMG preconditioners [13, 14, 15].

It is often advocated that the complexity of implementing these methods outweighs any gains made in efficiency over simpler techniques such as the Additive Schwarz procedures. It is hoped that this paper will demonstrate that this is not the case. It is possible to obtain good efficiencies while using only the local data structures used by the Schwarz procedures. A similar point was made in [16] using a Schur complement viewpoint. It was shown that global preconditioners can be designed from techniques which approximately solve the Schur complement system associated with interface variables. In this paper we extend this idea by using the Algebraic Recursive Multilevel Solvers (ARMS) framework.

The paper is organized as follows. Section 2 gives an overview of the sequential ARMS preconditioner. Distributed Schur complement and Additive Schwarz techniques, which serve as building blocks for the proposed parallel ARMS implementation, are outlined in Section 3. A more detailed description of the parallel ARMS design and implementation can be found in Section 4 followed by a few sets of numerical experiments, in Sections 5 and 6. Concluding remarks are presented in Section 7.

## 2. Sequential ARMS – basic notions

The multi-level ILU preconditioners developed in [1, 3, 4, 5, 17] exploit the property that a set of unknowns that are not coupled with each other can be eliminated simultaneously in Gaussian elimination. Such sets are termed “independent sets”, see e.g., [18]. In [4], the ILUM factorization described in [1] was generalized by resorting to “group-independent sets”. A group-independent set is a set of groups of unknowns such that there is no coupling between unknowns of any two different groups [4]. Unknowns within the same group may be coupled. This is illustrated in Figure 1. In the Algebraic Multigrid literature, certain types of group-independent sets are referred to as “aggregates” [19]. Simple methods for finding standard (point) and group (block) independent sets have been considered in [1, 4] and elsewhere.

In many existing forms of multilevel ILU factorizations [1, 2, 3, 9] the unknowns are reordered, listing the nodes associated with the group independent sets first, followed by the




Figure 1. Group independent sets.

other unknowns. After this reordering, the original system  $A_l x_l = b_l$  at the  $l$ -th level takes the following form

$$\begin{pmatrix} B_l & F_l \\ E_l & C_l \end{pmatrix} \begin{pmatrix} u_l \\ y_l \end{pmatrix} = \begin{pmatrix} f_l \\ g_l \end{pmatrix}. \tag{1}$$

At the zero-th level ( $l = 0$ ) the system  $A_l x_l = b_l$  is the original coefficient matrix of the linear system under consideration. The above reordered matrix is then approximately factored as

$$\begin{pmatrix} B_l & F_l \\ E_l & C_l \end{pmatrix} \approx \begin{pmatrix} L_l & 0 \\ E_l U_l^{-1} & I \end{pmatrix} \times \begin{pmatrix} U_l & L_l^{-1} F_l \\ 0 & A_{l+1} \end{pmatrix}, \tag{2}$$

where  $I$  is the identity matrix,  $L_l$  and  $U_l$  form the LU (or ILU) factors of  $B_l$ , and  $A_{l+1}$  is an approximation to the Schur complement with respect to  $C_l$ ,

$$A_{l+1} \approx C_l - (E_l U_l^{-1})(L_l^{-1} F_l). \tag{3}$$

During the factorization process, approximations to the matrices  $G_l \equiv E_l U_l^{-1}$  and  $W_l \equiv L_l^{-1} F_l$  are computed for obtaining the Schur complement (3) but these two matrices are discarded after  $A_{l+1}$  is computed. Typically it is inexpensive to solve linear systems with  $U_l$  and  $L_l$  since these arise from an ILU-type factorization. Therefore, observe that all we need for defining a preconditioning for  $A_l$  is to provide a way to solve the reduced system, i.e., the system associated with  $A_{l+1}$  obtained by eliminating the unknowns associated with the block  $B_l$ . It is here that different methods proposed vary.

Solving a linear system with the matrix in (2) requires the following three operations: a forward solve followed by a solve with the coarser level matrix  $A_{l+1}$ , and followed by a backward solve. The first and third operations move from one level to another and are similar to the restriction and prolongation operations in multigrid techniques. We refer to these as a downward and an upward operation, respectively. The descend operation leads to the system

$$A_{l+1} \times y_l = b_{l+1}, \quad b_{l+1} = g_l - E_l U_l^{-1} L_l^{-1} f_l. \tag{4}$$

Notice that  $y_l$  and  $x_{l+1}$  denote the same unknown. At any given step, the system at the next level, i.e., the system (4), is solved by an unspecified approximate or exact solution technique. The variations that arise are related to the ways in which this coarser level system is solved. In [6], we have implemented several different options and tested them. The simplest option available is referred to as VARMS [in analogy with the V-cycle multigrid]. If the current

level is not the last, VARMS continues to descend by using the level-structure, then solves the last-level system with GMRES-ILUT, and ascends back to the current level. Another option, not considered in this paper, is referred to as WARMS. At the difference with VARMS, WARMS iterates when solving the intermediate linear systems by using a few steps of GMRES with VARMS as a preconditioner. The last level system is again solved with GMRES-ILUT. For reference, we reproduce here the description of the VARMS recursive solution from [6].

**ALGORITHM 2.1.** *VARMS-solve( $A_l, b_l$ ) – Recursive Multi-Level Solution*

1. Solve  $L_l f'_l = f_l$
2. Descend, i.e., compute  $b_{l+1} := g_l - E_l U_l^{-1} f'_l$
3. If  $l = \text{last\_lev}$  then
4.     Solve  $A_{l+1} y_l = b_{l+1}$  using GMRES+ILU factors
5. Else
6.      $y_{m+1} = \text{VARMS-solve}(A_{l+1}, b_{l+1})$
7. Endif
8. Ascend, i.e., compute  $f''_l = f'_l - L_l^{-1} F_l y_l$
9. Back-Substitute  $u_l = U_l^{-1} f''_l$

### 3. Preconditioning of distributed sparse linear systems

The framework we adopt for solving large sparse linear systems on parallel platforms is that of a distributed sparse linear system. This viewpoint generalizes Domain Decomposition methods to irregularly structured sparse linear systems. A typical distributed system arises, for example, from a finite element discretization of a partial differential equation on a certain domain. To solve such systems on a distributed memory computer, it is common to partition the finite element mesh by a graph partitioner and assign a cluster of elements which represent a physical subdomain to each processor. Each processor then assembles only the local equations associated with the elements assigned to it. The general assumption is that each processor holds a set of equations (rows of the linear system) and a vector of the variables associated with these rows.

Figure 2 shows a ‘physical domain’ viewpoint of a sparse linear system. As is often done, we will distinguish between three types of unknowns: (1) Interior variables are those that are coupled only with local variables by the equations; (2) Interdomain interface variables are those coupled with non-local (external) variables as well as local variables; and (3) External interface variables are those variables in other processors which are coupled with local variables. The local equations can be represented as shown in Figure 3. Note that these equations are not contiguous in the original system. The matrix represented in the figure can be viewed as a reordered version of the equations associated with a local numbering of the equations/unknowns pairs.

As can be seen in Figure 3, the rows of the matrix assigned to a certain processor have been split into two parts: a *local* matrix  $A_i$  which acts on the local variables and an *interface* matrix  $X_i$  which acts on remote variables. These remote variables must be first received from other processor(s) before the matrix-vector product can be completed in these processors. Most data structures for distributed sparse matrices list the interface nodes separately, usually by ordering them after the interior nodes. This ‘local ordering’ of the data presents several advantages,




Figure 2. A local view of a distributed sparse matrix.




Figure 3. A partitioned sparse matrix.

including more efficient interprocessor communication, and reduced local indirect addressing during matrix-vector products. It should be noted that the use of block sparse row format can also yield a substantial reduction in indirect addressing. The zero blocks in the figure shown are due to the fact that local internal nodes are not coupled with external nodes.

Thus, each local vector of unknowns  $x_i$ , ( $i = 1, \dots, p$ ) is split in two parts: the subvector  $u_i$  of internal nodes followed by the subvector  $y_i$  of interdomain interface variables. The right-hand side  $b_i$  is conformally split in the subvectors  $f_i$  and  $g_i$ ,

$$x_i = \begin{pmatrix} u_i \\ y_i \end{pmatrix}; \quad b_i = \begin{pmatrix} f_i \\ g_i \end{pmatrix}. \quad (5)$$

The local matrix  $A_i$  residing in processor  $i$  as defined above is block-partitioned according to

this splitting, leading to

$$A_i = \left( \begin{array}{c|c} B_i & F_i \\ \hline E_i & C_i \end{array} \right). \quad (6)$$

With this, the local equations can be written as follows:

$$\begin{pmatrix} B_i & F_i \\ E_i & C_i \end{pmatrix} \begin{pmatrix} u_i \\ y_i \end{pmatrix} + \begin{pmatrix} 0 \\ \sum_{j \in N_i} E_{ij} y_j \end{pmatrix} = \begin{pmatrix} f_i \\ g_i \end{pmatrix}. \quad (7)$$

The term  $E_{ij}y_j$  is the contribution to the local equation from the neighboring subdomain number  $j$  and  $N_i$  is the set of subdomains that are neighbors to subdomain  $i$ . The sum of these contributions, seen on the left side of (7), is the result of multiplying a certain matrix by the external interface variables. It is clear that the result of this product will affect only the interdomain interface variables as is indicated by the zero in the upper part of the second term in the left-hand side of (7). For practical implementations, the subvectors of external interface variables are grouped into one vector called  $y_{i,ext}$  and the notation

$$\sum_{j \in N_i} E_{ij} y_j \equiv X_i y_{i,ext}$$

will be used to denote the contributions from external variables to the local system (7). In effect, this represents a local ordering of external variables to write these contributions in a compact matrix form. With this notation, the left-hand side of (7) becomes

$$w_i = A_i x_i + X_{i,ext} y_{i,ext}. \quad (8)$$

Note that  $w_i$  is also the local part of the matrix-vector product  $Ax$  in which  $x$  is a vector which has the local vector components  $x_i$ .

### 3.1. Additive Schwarz preconditioning

Preconditioners for distributed sparse linear systems are best designed from the local structure described above. Additive and (variants of) multiplicative Schwarz procedures are the simplest preconditioners available. Additive Schwarz procedures, update the local solution by the vector obtained from solving the linear system formed by the local matrix and the local residual. The exchange of data is done through the computation of the residual. In simple terms, the Additive Schwarz preconditioners can be stated as follows:

**ALGORITHM 3.1.** *Additive Schwarz*

1. Update local residual  $r_i = (b - Ax)_i$
2. Solve  $A_i \delta_i = r_i$
3. Update local solution  $x_i = x_i + \delta_i$

This loop is executed on each processor simultaneously. Exchange of information takes place in Line 1, when the (global) residual is updated. Note that the residual is “updated” in that only the  $y$ -part of the right-hand side is changed. The local systems  $A_i \delta_i = r_i$  can be solved in three ways: (1) By a (sparse) direct solver, (2) by using a standard preconditioned Krylov solver, or (3) by performing a backward-forward solution associated with an accurate ILU (e.g., ILUT) preconditioner. Experiments show that option (3) or option (2) with only a very small number of inner steps (e.g., 5) is quite effective.

### 3.2. Schur complement techniques

Schur complement techniques refer to methods which iterate on the interdomain unknowns only, implicitly using internal unknowns as intermediate variables. These techniques are at the basis of what will be described in the next sections. Schur complement systems are derived by eliminating the variable  $u_i$  from the system (7). Extracting from the first equation  $u_i = B_i^{-1}(f_i - F_i y_i)$  yields, upon substitution in the second equation,

$$S_i y_i + \sum_{j \in N_i} E_{ij} y_j = g_i - E_i B_i^{-1} f_i \equiv g'_i, \quad (9)$$

where  $S_i$  is the “local” Schur complement

$$S_i = C_i - E_i B_i^{-1} F_i. \quad (10)$$

The equations (9) for all subdomains  $i$  ( $i = 1, \dots, p$ ) constitute a global system of equations involving only the interface unknown vectors  $y_i$ . This global reduced system has a natural block structure related to the interface points in each subdomain:

$$\begin{pmatrix} S_1 & E_{12} & \dots & E_{1p} \\ E_{21} & S_2 & \dots & E_{2p} \\ \vdots & & \ddots & \vdots \\ E_{p1} & E_{p-1,2} & \dots & S_p \end{pmatrix} \begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_p \end{pmatrix} = \begin{pmatrix} g'_1 \\ g'_2 \\ \vdots \\ g'_p \end{pmatrix}. \quad (11)$$

The diagonal blocks in this system, the matrices  $S_i$ , are dense in general. The off-diagonal blocks  $E_{ij}$ , which are identical with those involved in the system (7), are sparse.

The system (11) can be written as

$$S y = g',$$

where  $y$  consists of all interface variables  $y_1, y_2, \dots, y_p$  stacked into a long vector. The matrix  $S$  is the “global” Schur complement matrix. An idea proposed in [16] is to exploit methods that *approximately solve the reduced system* (11) to develop preconditioners for the original (global) distributed system. The resulting “global” preconditioner is said to be “induced” from the technique used for the Schur complement. Once the global Schur complement system (9) is (approximately) solved, each processor will compute the  $u$ -part of the solution vector (see (5)) by solving the system  $B_i u_i = f_i - E_i y_i$  obtained by substitution from (7). In summary, a Schur complement iteration may be expressed by the following algorithm:

**ALGORITHM 3.2.** *Schur Complement Iteration — Template*

1. *Forward: compute local right-hand sides  $g'_i = g_i - E_i B_i^{-1} f_i$*
2. *Solve global Schur complement system  $S y = g'$*
3. *Backward: substitute to obtain  $u_i$ , i.e., solve  $B_i u_i = f_i - E_i y_i$*

Here, the similarity with ARMS is worth mentioning. The first step of every Schur complement algorithm considered in this paper is to form the new right-hand side  $g'_i = g_i - E_i B_i^{-1} f_i$ . This step is identical to the forward operation used for ARMS and corresponds to solving with the L factor in a block LU factorization (see [16]). The second step requires that we (approximately) solve the Schur complement system (9) in some way. At this stage, the  $y$  variables are approximated. Finally, the third operation is the backward step, which computes the rest of the unknowns by substitution.

We now consider step two, solution of system (9), in more detail. For convenience, (9) is rewritten as a preconditioned system with the diagonal blocks:

$$y_i + S_i^{-1} \sum_{j \in N_i} E_{ij} y_j = S_i^{-1} [g_i - E_i B_i^{-1} f_i]. \quad (12)$$

This can be viewed as a block-Jacobi preconditioned version of the Schur complement system (11). This global system can be solved by a GMRES-like accelerator, requiring a solve with  $S_i$  at each step.

In [16], two techniques were examined for solving linear systems with  $S_i$ . The first one is based on the observation that an LU (or ILU) factorization of  $S_i$  can be easily extracted from an LU (or ILU) factorization of  $A_i$ . Specifically, if  $A_i$  has the form (6) and it is factored as  $A_i = L_i U_i$ , where

$$L_i = \begin{pmatrix} L_{B_i} & 0 \\ E_i U_{B_i}^{-1} & L_{S_i} \end{pmatrix} \quad \text{and} \quad U_i = \begin{pmatrix} U_{B_i} & L_{B_i}^{-1} F_i \\ 0 & U_{S_i} \end{pmatrix}.$$

then,  $L_{S_i} U_{S_i}$  is equal to the Schur complement  $S_i$  associated with the partitioning (6). In other words, *an ILU factorization for the Schur complement is the trace of the global ILU factorization on the unknowns associated with the Schur complement*. For a local Schur complement, the ILU factorization obtained in this manner leads to an approximation  $\tilde{S}_i$  of the local Schur complement  $S_i$ . The second method examined in [16] was based on using approximate inverse techniques for computing an approximation to the local Schur complement  $S_i$ .

#### 4. Parallel implementation of ARMS (pARMS)

In order to extend ARMS to a parallel framework, it is necessary to begin by generalizing the notion of group-independent sets into Distributed Group-Independent sets (DGIS). A DGIS is a set of subsets (groups) from all processors such that the unknowns of different groups (within and across processors) are not coupled. It is clear that there is already one such set available from the partitioning which is the set consisting of the internal nodes, subdomain by subdomain. Distributed group-independent sets are easily obtained by further subdividing these sets of interior nodes. Doing so will yield another category of points: local interface points. An illustration is shown in Figure 4 for a simple case with two subdomains. The thick line in the middle contains no vertices – it only illustrates the separation of the two domains. The partitioning, which is vertex-based, is the same as the one seen in the previous section. Once the system is distributed among processors, using a graph partitioner, an independent set reordering is applied locally. One variation is that the independent set ordering does not disturb, i.e., it does not reorder, the interdomain interface points. Instead, independent sets are found among the interior points only.

With this viewpoint, the local matrix can be represented as in Figure 5, in which IS represents the points in distributed group-independent sets, I1 represents the set of local interface points, while I2 represents the set of interdomain interface points. In parallel ARMS, the points I1 and I2 are all treated as interface points, in the sense that the Schur complement system considered may include both sets of points. When both types of points are included, we will refer to this system as the *expanded Schur complement*. It is clear that, when dealing with two types of






Figure 4. A two-level group-independent set.




Figure 5. Matrix associated with 2-level partitioning

interface points, some differences will arise. For example, I1-points require no communication when solving the linear systems.

There are three distinct ways of using ARMS in a parallel environment. The simplest is to use an Additive Schwarz procedure in which the local solver uses the ARMS preconditioner. The second approach relies on a Schur complement-type technique in which, as was just explained, the Schur complement relates to equations associated with (local and inter-processor) interface points. This second approach yields a whole set of algorithms which arise from the different ways the Schur complement system is solved. We can solve the Schur complement system by a block-Jacobi procedure, and this would be similar to the work in [16]. In these two approaches, the ARMS reordering is applied locally only. If the system to solve is very large, the Schur complement system becomes itself quite costly to solve and it may be necessary to extend the ARMS reordering for the interdomain interface variables. This constitutes the third approach which is not considered in this paper. In the rest of this section, we present details of the expanded Schur complement system and describe the preconditionings acting on it.

#### 4.1. Global expanded Schur complement system

In this subsection, the local processor index  $i$ ,  $i = (1, \dots, p)$ , will be omitted whenever it is clear that the related operations are local to processor  $i$ . The partitioning of the interior unknowns  $u$  into group-independent sets defines their splitting into interior points  $\tilde{u}$  and local interface points  $\hat{u}$ . The local system (7) can be permuted accordingly into a block form in which the points  $\tilde{u}$  are labeled first followed by the local interface points  $\hat{u}$  and interdomain interface points  $y$  (see also Figure 4):

$$\begin{pmatrix} \tilde{B} & \hat{F} & F' \\ \hat{E} & \hat{C} & F'' \\ E' & E'' & C \end{pmatrix} \begin{pmatrix} \tilde{u} \\ \hat{u} \\ y \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ \sum_{j \in N} E_j y_j \end{pmatrix} = \begin{pmatrix} \tilde{f} \\ \hat{f} \\ g \end{pmatrix}. \quad (13)$$

In system (13), the matrix  $\tilde{B}$  is block-diagonal. By eliminating the unknowns  $\tilde{u}$  from (13) in each processor, the *global expanded* Schur complement system, which corresponds to the remaining unknowns  $z = (\hat{u}, y)$ , is formed.

$$S z + \begin{pmatrix} 0 \\ \sum_{j \in N} E_j y_j \end{pmatrix} = \begin{pmatrix} \hat{f} - \hat{E} \tilde{B}^{-1} \tilde{f} \\ g - E' \tilde{B}^{-1} \tilde{f} \end{pmatrix} \equiv g', \quad (14)$$

where  $S$  is the ‘‘local’’ Schur complement, which corresponds to the local and interdomain interface nodes:

$$S = \begin{pmatrix} \hat{C} & F'' \\ E'' & C \end{pmatrix} - \begin{pmatrix} \hat{E} \tilde{B}^{-1} \hat{F} \\ E' \tilde{B}^{-1} \hat{F} \end{pmatrix}.$$

The structure of this expanded Schur complement system is similar to the one in equation (11) and also has a natural block structure related to the interface points in each subdomain.

#### 4.2. Additive Schwarz for the Schur complement

This consists of applying the Schur complement iteration (Algorithm 3.2) using the Schur complement system derived from pARMS. The system (14), the expanded version of (9), can be rewritten similarly to (12), as a preconditioned system with the diagonal blocks

$$z + S^{-1} \begin{pmatrix} 0 \\ \sum_{j \in N} E_j y_j \end{pmatrix} = S^{-1} g', \quad (15)$$

This is simply a block-Jacobi preconditioned version of the global expanded Schur complement system (14). In [16], incomplete LU factorizations and approximate inverses were used to approximate the local  $S^{-1}$ . Here, we apply a local ARMS procedure to obtain similar approximations. In particular, when  $S$  is the last level  $k$  of the local ARMS recursion, i.e.,  $S = A_k$ , a factorization can be defined as follows. Let  $A_{k-1}$  be block factored as in (6) and followed by an (incomplete) LU factorization of the last matrix  $A_k = L_k U_k$ . Taken separately, this part of the ARMS factorization can be used for the preconditioning of the Schur complement system. For LU factorizations, this result has been established in [8]. In the case of the ARMS preconditioner, it can be easily extended when local ARMS recursively factors  $S$ . For a local Schur complement, an ARMS factorization obtained in this manner leads to an approximation  $\tilde{S}$  of the local Schur complement  $S$ . Instead of the exact Schur complement system (14), or equivalently (15), the following approximate (local) Schur complement system, derived from (15), can be considered in each processor:

$$z + \tilde{S}^{-1} \begin{pmatrix} 0 \\ \sum_{j \in N} E_j y_j \end{pmatrix} = \tilde{S}^{-1} g'. \quad (16)$$

An algorithm is presented in [16] that applies, in each processor, the global approximate Schur LU preconditioner to a block vector  $(f, g)^T$  to obtain the solution  $(u, y)^T$ . The algorithm uses a small number of GMRES iterations without restarting to approximate the solution of the local part of the Schur complement system (16).

#### 4.3. Global incomplete LU preconditioning

In the previous subsection, the Schur complement system is solved by a Krylov subspace accelerator with block-Jacobi as a preconditioner. More accurate solutions to the Schur complement system may lead to a better convergence of the outer iteration. The distributed ILU(0) preconditioner may be a good alternative to a block-Jacobi preconditioner or to an SSOR preconditioner described [20]. We may use this approach as well to solve the Schur complement system.

The Schur complement system considered here is again the expanded system which contains local and interdomain interface points (see Figure 4). Interior points can be eliminated first, simultaneously in each processor. Once this is done, the interdomain interface points can be eliminated *in a certain order*. This is no different from standard distributed ILU(0) [8], except that it is applied to the distributed Schur system instead of the original distributed system. The global ordering used is based on a multicoloring of the domains, as this yields good parallelism [8].

As was already mentioned the distributed ILU(0) preconditioner is applied to the *local interface* and *interdomain interface* points. This technique is referred to as Schur Global ILU(0) preconditioner, to distinguish it from the distributed ILU(0) applied to the original distributed system. In what follows *colors*[ $i$ ] is the color of processor number  $i$ , *mycolor* is the color of the current processor, *ncolor* is the number of colors.

#### ALGORITHM 4.1. Global ILU(0) factorization

1. Perform ARMS elimination on the rows associated with local group-independent sets.
2. Perform ILU(0) on the local interface rows.
3. /\* Perform ILU(0) on the interdomain interface rows: \*/

```

4.   for( $i = 0$ ;  $i < ncolor$ ;  $i ++$ ) {
5.       if ( $mycolor == i$ ) {
6.           Perform the  $ILU(0)$  factorization for the interdomain interface
              rows with pivots from the local interface rows completed in
              step 2;
7.           Send the completed interdomain interface rows to adjacent
              processors  $j$ , with  $colors[j] > mycolor$ .
              }
8.       else if ( $mycolor > i$ ) {
9.           Receive the factored rows from the adjacent processors
               $j$  with  $colors[j] == i$ ;
10.          Perform the  $ILU(0)$  factorization with pivots received from
              the external processors in step 7.
              }
          }
    }

```

In the above algorithm, steps 1 and 2 can be performed simultaneously in each processor.

The preconditioned Krylov subspace algorithm requires a forward and backward sweep at each step. The distributed forward/backward solution based on this factorization is sketched next.

**ALGORITHM 4.2.** *Distributed forward and backward sweep*

```

1.   /* Forward solve: */
2.       Perform the forward solve for the local interface nodes.
3.       for( $i = 0$ ;  $i < ncolor$ ;  $i ++$ ) {
4.           if ( $mycolor == i$ ) {
5.               Perform the forward solve for the interdomain interface nodes;
6.               Send the updated values of interdomain nodes to the adjacent
              processors  $j$ , with  $colors[j] > i$ .
              }
7.           else if ( $mycolor > i$ ) {
8.               Receive the updated values from the adjacent processors
               $j$ , with  $colors[j] == i$ .
              }
          }
9.   /* Backward solve: */
10.      for( $i = ncolor - 1$ ;  $i >= 0$ ;  $i --$ ) {
11.          if ( $mycolor == i$ ) {
12.              Perform the backward solve for the interdomain interface nodes;
13.              Send the updated values of interdomain nodes to the adjacent
              processors  $j$ , with  $colors[j] < i$ .
              }
14.          else if ( $mycolor < i$ ) {
15.              Receive the updated values from the adjacent processors
               $j$ , with  $colors[j] == i$ .
              }
          }
    }

```

16. *Perform the backward solve for the local interface nodes.*

The above algorithm only shows the forward/backward sweep in the last Schur complement. As in the ILU(0) factorization, the local interface nodes do not depend on the nodes from the external processors and can be computed in parallel in steps 2 and 16. In the forward solve, the solution of the local interface nodes is followed by an exchange of data and the solution on the interdomain interface. The backward solve works in reverse in that the interdomain interface nodes are first computed, then they are sent to adjacent processors, and the local interface update follows.

## 5. Numerical Experiments

In this section, we present a comparison of the preconditioners available in the pARMS framework. Scalability experiments, in which the problem size and number of processors grow, are first reported for two- and three-dimensional PDE problems on regular grids. A performance comparison with a direct solver is included in these tests. We then report on a few tests with fixed-size irregularly structured problems arising in magnetohydrodynamics.

### 5.1. Naming conventions for the methods tested

For the sake of convenience, let us introduce some notation for the preconditioning options considered in pARMS. In general, the method names are of the form `add_X Y` or `sch_X Y`, where the prefixes `add` and `sch` indicate whether the preconditioner has local operations only, i.e., it is of Additive Schwarz type, or whether it requires external data communications, i.e., it is of Schur complement type, respectively. The letter `X` may stand for either an incomplete LU factorization, such as ILUT or ILU(k) [8], or for ARMS. The global preconditioners of the `sch` type may be also used with a global ILU(0) preconditioner `gilu0` applied to the expanded Schur complement system. This preconditioning option `sch_gilu0` resembles the one in [16], but it acts on the expanded Schur complement system (13) instead of (9), and it uses an ILU(0) factorization instead of a block Jacobi preconditioner. A preconditioner name may be followed by `Y`, which stands for `no_its` and/or `ovp`, meaning that the preconditioner is applied without inner (i.e., intra-domain) iterations and/or has overlapping, respectively. As is well-known [21, 22, 23, 24, 25] overlapping improves the efficiency and scalability of Additive Schwarz preconditioners. Many ways to overlap nodes have been considered in parallel environments. In particular, Restrictive Additive Schwarz [26] showed good performance. In the experiments, we consider a simple overlapping technique, in which the contributions of the overlapped nodes are averaged upon their exchange among neighboring domains. This type of overlapping has been used by many authors, see, for example, [27, 28, 29]. A few additional details are provided below for each of the preconditioners considered.

`add_ilut`. Additive Schwarz procedure, with or without overlapping, in which ILUT is used as a preconditioner for solving the local systems. These systems can be solved with a given number of GMRES inner iterations or by just applying the preconditioner.

`add_iluk`. Similar to `add_ilut` but uses ILU(k) as a preconditioner instead of ILUT.

`add_arms`. Similar to `add_ilut` but uses ARMS as a preconditioner for local systems.

`sch_ilut`. This method is the same as the one in [16]. It consists of solving the (global) Schur complement system, associated with the interdomain interface variables, using a block-Jacobi preconditioner. The ILUT factorization for the local Schur complement is obtained as a ‘trace’ of the ILUT factorization on the interface variables. These local ILUT factorizations are then used to precondition the global Schur complement system in a block-Jacobi iteration. For details see [16].

`sch_iluk`. This is the same as `sch_ilut` except that ILU(k) is used instead of ILUT.

`sch_arms`. This is the same as `sch_ilut` except that ARMS is used instead of ILUT.

`sch_sgs`. Symmetric block Gauss-Seidel preconditioning is used in solving the global Schur complement system instead of a block-Jacobi preconditioning as described in [16]. This preconditioning requires a global ordering, which is provided by multicoloring.

`sch_gilu0`. This method differs from the other Schur complement methods in that it is based on approximately solving the expanded Schur complement with a global ILU(0)-preconditioned GMRES. The ILU(0) preconditioning requires a global order (referred to as a schedule in [30]) in which to process the nodes. A global multicoloring of the domains is used for this purpose as is often done with global ILU(0).

## 5.2. Scaled problem: Two-dimensional case

Consider the elliptic partial differential equation

$$-\Delta u + 100 \frac{\partial}{\partial x} (e^{xy} u) + 100 \frac{\partial}{\partial y} (e^{-xy} u) - 10u = f \quad (17)$$

on a square region with Dirichlet boundary conditions, discretized with a five-point centered finite-difference scheme on a  $n_x \times n_y$  grid, excluding boundary points. The mesh is mapped to a virtual  $p_x \times p_y$  grid of processors, such that a subrectangle of  $r_x = n_x/p_x$  points in the  $x$  direction and  $r_y = n_y/p_y$  points in the  $y$  direction is mapped to a processor. In the following experiments, the mesh-size in each processor is kept constant at  $r_x = 100, r_y = 100$ . As the number of processors increases, the problem size increases proportionally. For example, on 4 processors the mesh is  $(2r_x, 2r_y)$  leading to a problem size of 40,000 and on 64 processors the mesh is  $(8r_x, 8r_y)$  leading to a problem size of 640,000. The resulting problems become harder as the number of processors grows. In a perfectly scalable situation, the final execution time for solving the problem with size  $n = 640,000$  on 64 PEs should be identical with the time for solving the problem of size  $n = 40,000$  on 4 PEs. Note that in order to maintain the aspect ratio of the physical domain, we need to consider *square* processor grids of increasing size.

The results have been obtained on an SGI Origin 3800. The residual norm reduction of  $10^{-6}$  was achieved by flexible GMRES (FGMRES) [8] with subspace size of 100. The fill-in parameter `lfil` is taken as 60 for all the levels of recursion along with the dropping tolerance of  $10^{-4}$ . The group-independent set size is 20. Timing and iteration results for the fixed local subproblem sizes of  $100 \times 100$  are presented in Figures 6 and 7 (top and bottom, respectively). As shown in Figure 6, the total solution times for the Additive Schwarz preconditioners with five inner iterations are larger than for the corresponding preconditioners without inner

Table I. Comparison of preconditioners for 3D PDE problem

Preconditioner	NP = 8		NP = 27		NP = 64	
	iter	time	iter	time	iter	time
add_arms no its	24	15.64	32	19.31	41	25.72
add_arms no its ovp	13	13.58	20	18.05	26	22.37
sch_arms	10	15.10	14	19.35	20	30.31
sch_gilu0	8	13.9	11	18.26	15	31.35
sch_gilu0 no its	15	12.98	22	16.72	29	20.93
sch_sgs	8	15.68	11	20.73	14	28.83
sch_sgs no its	13	13.72	19	17.12	24	22.00

iteration. For Additive Schwarz, the presence of inner iterations does not reduce the number of outer iterations in a significant way. Additive Schwarz with ARMS as a local preconditioner uses less time than Additive Schwarz with ILUT as a local preconditioner although their corresponding iteration numbers are close. The reason is that Additive Schwarz with ARMS requires less memory in both construction and solution. Figures 6 and 7 confirm the advantages of using overlapping in Additive Schwarz.

Figure 7 shows a comparison of two-level preconditioners with various Additive Schwarz procedures. It is noticeable that `sch_gilu0` and `sch_sgs` outperform the other methods for this test case. The performance of `sch_ilut` has also been tested. However, it is not shown in the Figures since this method was not found to be faster than `sch_gilu0` or `sch_sgs` and its advantages over `add_ilut` have been shown in our earlier work [16]. Note a rather interesting spread in the number of iterations for the various methods. It is also remarkable that `sch_sgs` and `sch_gilu0` require an almost identical number of steps to converge, both when inner iterations are performed and when no inner iterations are performed.

### 5.3. Scaled problem: Three-dimensional case

We solve an analogue of (17), on a three-dimensional cube. The PDE is expressed in an identical way as in (17) but the Laplacean is three-dimensional, i.e., there is no convective term along the  $z$  direction. The PDE is set on a cube with Dirichlet boundary conditions, and it is discretized with a seven-point centered finite-difference scheme on a  $n_x \times n_y \times n_z$  grid, excluding boundary points.

The experiments have been performed on 8, 27, and 64 processors of an IBM SP with the same parameters as in two-dimensional case. In Table I, the columns `iter` and `time` contain iterations numbers and total times, respectively. It can be seen that, for this problem type, `sch_gilu0 no its` outperforms other methods mainly because it requires less computational work.

### 5.4. Comparison of *pARMS* and *PSPASES*

We now compare *pARMS* with a parallel sparse direct solver. Among the parallel solvers available to us, we selected to use *PSPASES* (Scalable Parallel Direct Solver Library for Sparse

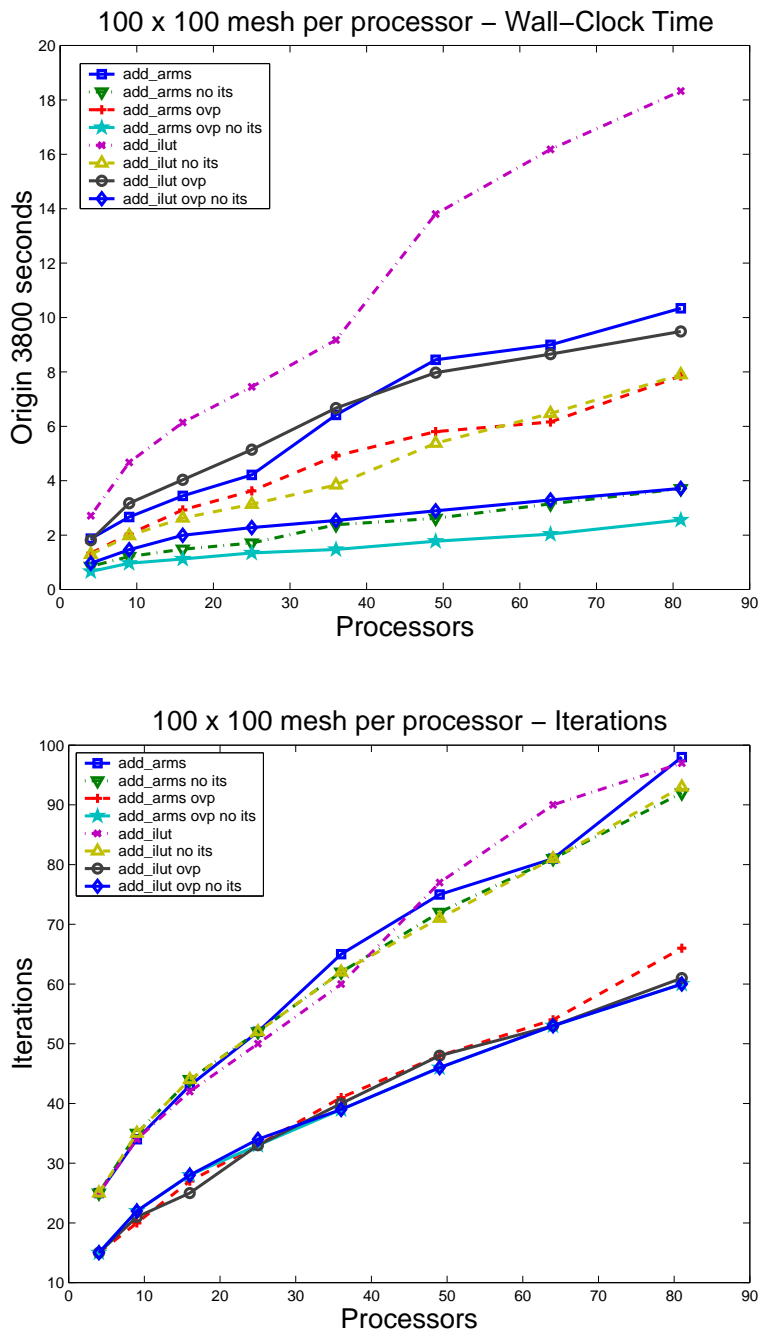


Figure 6. Solution times (top) and iterations (bottom) for a 2D PDE problem with the fixed subproblem size using variations of Additive Schwarz preconditioner.



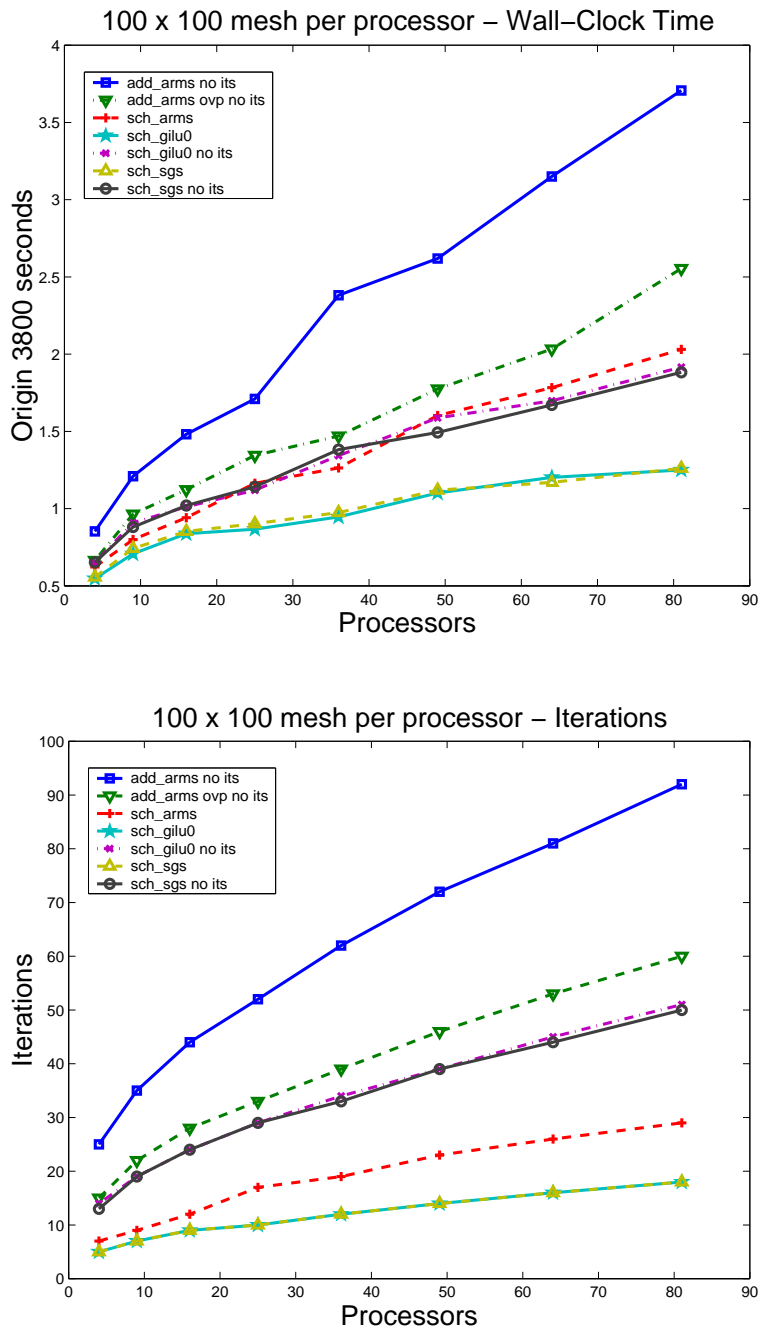


Figure 7. Solution times (top) and iterations (bottom) for a 2D PDE problem with the fixed subproblem size using four different preconditioners.

Table II. Comparison of pARMS and PSPASES for a 2D problem

NP	sch_gilu0			PSPASES		
	TTIME	CTIME	STIME	TTIME	FTIME	STIME
4	0.78	0.47	0.29	0.67	0.63	0.045
8	0.99	0.56	0.41	0.75	0.70	0.05
16	1.20	0.55	0.64	1.14	1.07	0.07
32	1.74	0.69	1.04	1.77	1.68	0.09
64	2.97	0.79	2.17	3.19	3.03	0.16

Table III. Comparison of pARMS and PSPASES on a 3D problem

NP	sch_gilu0			PSPASES		
	TTIME	CTIME	STIME	TTIME	FTIME	STIME
8	10.70	5.44	5.24	86.17	85.50	0.67
64	27.78	7.92	19.85	*	*	*

Symmetric Positive Definite Linear Systems) [31]. Because PSPASES is a static code in the sense that it does a symbolic factorization prior to the numerical factorization, it is extremely efficient. On the other hand the code is restricted to solving Symmetric Positive Definite matrices. Another restriction is that the number of processors should be a power of 2. For the tests, we use simply the Poisson equation

$$-\Delta u = f \quad (18)$$

on two-dimensional rectangular region with Dirichlet boundary conditions. The problem is discretized as in the previous examples, with centered finite differences. In the first experiment, we solve a two-dimensional Poisson equation and select the mesh size on each processor to be  $100 \times 100$ . PSPASES has been used with the default parameters supplied with the package distribution. In Table II, TTIME, CTIME, FTIME, and STIME stand for total time, preconditioner construction time, factorization time, and solution time respectively. All times are in seconds on the SGI Origin 3800. As shown in Table II, for two-dimensional problem and small numbers of processors (fewer than 64), PSPASES is better than pARMS. On 64 processors, pARMS is slightly faster than PSPASES. In the next test, we solve the Poisson equation on a three-dimensional cube with Dirichlet boundary conditions. The problem size is  $30 \times 30 \times 30$  per processor. The timing results are presented in Table III, in which the columns have the same meaning as in Table II and an asterisk indicates that the direct solver failed due to insufficient memory. For this problem, sch\_gilu0 is much faster than PSPASES, as Table III shows. This comparison gives only a rough idea on the relative performance of pARMS and PSPASES since pARMS is not yet an optimized code. In addition, the pARMS iterative solution of the system does not take advantage of the problem symmetry. Other direct solution codes we have tried were either too slow or not ready for a distributed parallel environment.

### 5.5. A problem issued from magnetohydrodynamic flow

The linear system solved in this section is a fairly hard problem which arises from Magnetohydrodynamic (MHD) flows. The flow equations are represented as coupled Maxwell's and the Navier-Stokes equations. The conservative magnetohydrodynamic system is modeled by the Maxwell equations, written as:

$$\frac{\partial \mathbf{B}}{\partial t} - \nabla \times (\mathbf{u} \times \mathbf{B}) + \eta \nabla \times (\nabla \times \mathbf{B}) + \nabla q = 0 \quad (19)$$

$$\nabla \cdot \mathbf{B} = 0, \quad (20)$$

where  $\eta$ ,  $\mathbf{B}$ ,  $\mathbf{u}$  and  $q$  are, respectively, the magnetic diffusivity coefficient, magnetic induction field, velocity field, and the scalar Lagrange multiplier for the magnetic-free divergence constraint. In fully coupled magnetohydrodynamics, this system is solved along with the incompressible Navier-Stokes equations

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} - \nu \nabla^2 \mathbf{u} + \nabla p = \mathbf{f} \quad (21)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (22)$$

where  $p$ ,  $\nu$  and  $\mathbf{f}$  are, respectively, pressure, kinematic viscosity, and body force. The coupling between the two systems is through the body force  $\mathbf{f} = \frac{1}{\mu}(\nabla \times \mathbf{B}) \times \mathbf{B}$ , which represents the Lorentz (Laplace) force due to the interaction between the current density  $\mathbf{j} = \frac{1}{\mu}(\nabla \times \mathbf{B})$  and the magnetic field, where  $\mu$  is the magnetic permeability.

It is uncommon to solve the fully coupled problem described by the equations (19-22) along with their coupling via the body forces, because this usually requires an excessive amount of memory. Instead, segregated approaches are often applied which alternatively solve the two coupled problems until a certain convergence criterion is satisfied. For time-dependent problems, these coupling iterations are embedded into the time-stepping procedure. For a few details on this problem, its discretization, and the segregated solution procedure see [32].

In the tests that follow, we solve problems which arise from the Maxwell equations only. In order to do this, a pre-set periodic induction field  $\mathbf{u}$  is used in Maxwell's equation (19). The physical region is the three-dimensional unit cube  $[-1, 1]^3$  and the discretization uses a Galerkin-Least-Squares discretization. The magnetic diffusivity coefficient is  $\eta = 1$ . Two systems were generated. The first one (denoted by MHD1) has  $n = 485,597$  unknowns and 24,233,141 nonzero elements. As was mentioned above the function  $q$  in (19) corresponds to Lagrange multipliers, which arise from imposing the magnetic-free divergence constraint. Its gradient should be zero at steady-state. The second system is obtained by simply imposing Dirichlet boundary conditions for  $q$ . This results in a slightly smaller linear system of size  $n = 470,596$  equations and  $nnz = 23,784,208$  nonzeros. As it turns out, the second system is slightly easier to solve by iterative methods than the first one. Though the actual right-hand side was supplied, we preferred to use an artificially generated one in order to check the accuracy of the process. A random initial guess was taken. Little difference in performance was seen when the actual right-hand and a zero vector initial guess were used instead.

The following input parameters have been chosen to solve this problem: the residual is to be reduced by  $10^6$ , the FGMRES restart value is 60, and at most five inner iterations may be allowed at each preconditioner application to attain 0.01 accuracy of the residual norm reduction. The parameters for ARMS are as follows: the number of levels, the independent-set

block size, fill-in factors, and dropping tolerances for the intermediate and the last levels are 5, 5000, 60, 60,  $10^{-5}$ , and  $10^{-4}$ , respectively. Large block size and fill-in factors are prompted by the hardness of the problem. The resulting total number of nonzero memory locations required for the preconditioner was about 2.5 times the number of nonzeros in the original matrix. We observed that all the methods without inner iterations experienced stagnation for the MHD1 problem. Additive Schwarz (`add_arms no its`) with or without overlap does not converge for any number of processors while the Schur global ILU(0) (`sch_gilu0 no its`) stagnates when executed on more than nine processors. On four and nine processors, `sch_gilu0 no its` converges in 188 and 177 iterations, respectively. On an IBM SP, this amounts to 2,223.43 and 1,076.27 seconds, respectively. This is faster than 2,372.44 and 1,240.23 seconds when five inner iterations are applied and the number of outer iterations decreases to 119 and 109 on four and nine processors, respectively. The benefits of iterating on the global Schur complement system are clear since the Schur complement-based preconditioners converge for all the processor numbers tested as indicated in Figure 8, which shows the timing results (top) and outer iteration numbers (bottom). This positive effect can be explained by the fact that the Schur complement system is computed with good accuracy.

Comparing the two preconditioners in Figure 8 (bottom), we observe that `sch_gilu0` always requires fewer than `sch_arms` outer iterations and shows no iteration increase with the increase in processor numbers. The construction of `sch_arms`, however, is somewhat cheaper (cf. the “star” and “cross” curves in Figure 8, top) because it involves no interprocessor communication. The `sch_gilu0` application time advantage, seen for smaller processor numbers, vanishes as more processors are used (cf. the “rhombus” and “triangle” curves) since the cost of `sch_gilu0` increases compared with that of `sch_arms` as the number of processors increases, whereas the difference in numbers of iterations stays almost constant, for the MHD1 problem.

The MHD2 problem is easier to solve. In fact, this is a good example where iterating on the global Schur Complement is not mandatory. The convergence may be achieved with a simple Additive Schwarz-type preconditioner (`add_arms no its`) without overlapping, which, in this case, is quite competitive with respect to time (Figure 9, top). This is in spite of the rather high iteration counts for `add_arms no its`, relative to the other methods tested (Figure 9, bottom). All the Schur complement techniques produce fewer iterations to convergence. Applying `sch_gilu0` without inner iteration decreases solution times and makes `sch_gilu0 no its` faster than `add_arms not its` up to 32 processors.

In summary, the purpose of presenting these two examples of realistic problems is to show the flexibility of the developed pARMS framework. The difficulty and size of the problem at hand can call for a more powerful preconditioner, which can be readily provided by increasing fill-in and allowing more iterations on the global Schur complement variables. For easier problems (e.g., MHD2), Additive Schwarz methods or inexpensive variations of Schur complement techniques may be adequate and are provided as particular cases within the pARMS package.

## 6. A comparison with other parallel preconditioners

We have found only a limited selection of packages which include parallel preconditioners for general sparse linear systems. Among these we have selected the implementations available in the *hypr* package [13] from Lawrence Livermore National Laboratory for a comparison. This

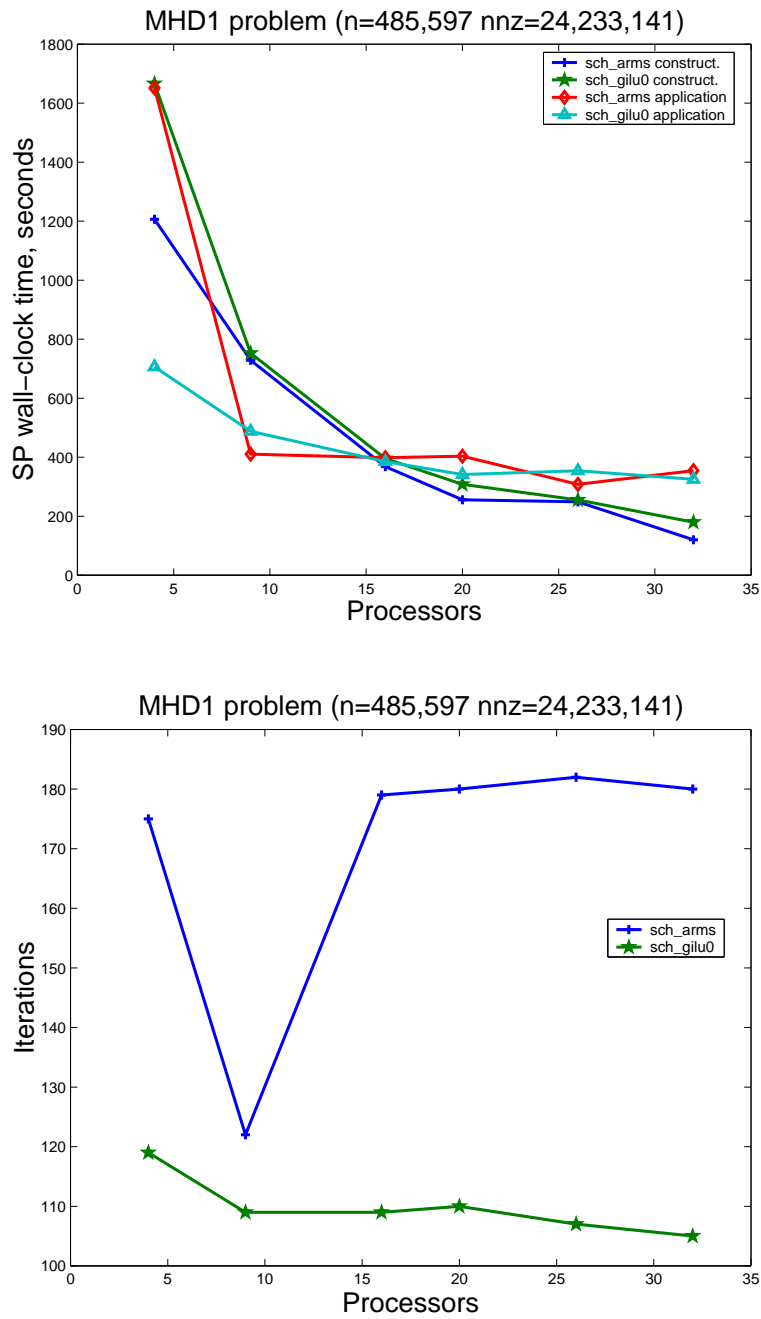


Figure 8. Timing results (top) and outer iterations (bottom) for the (fixed-size) MHD1 problem.

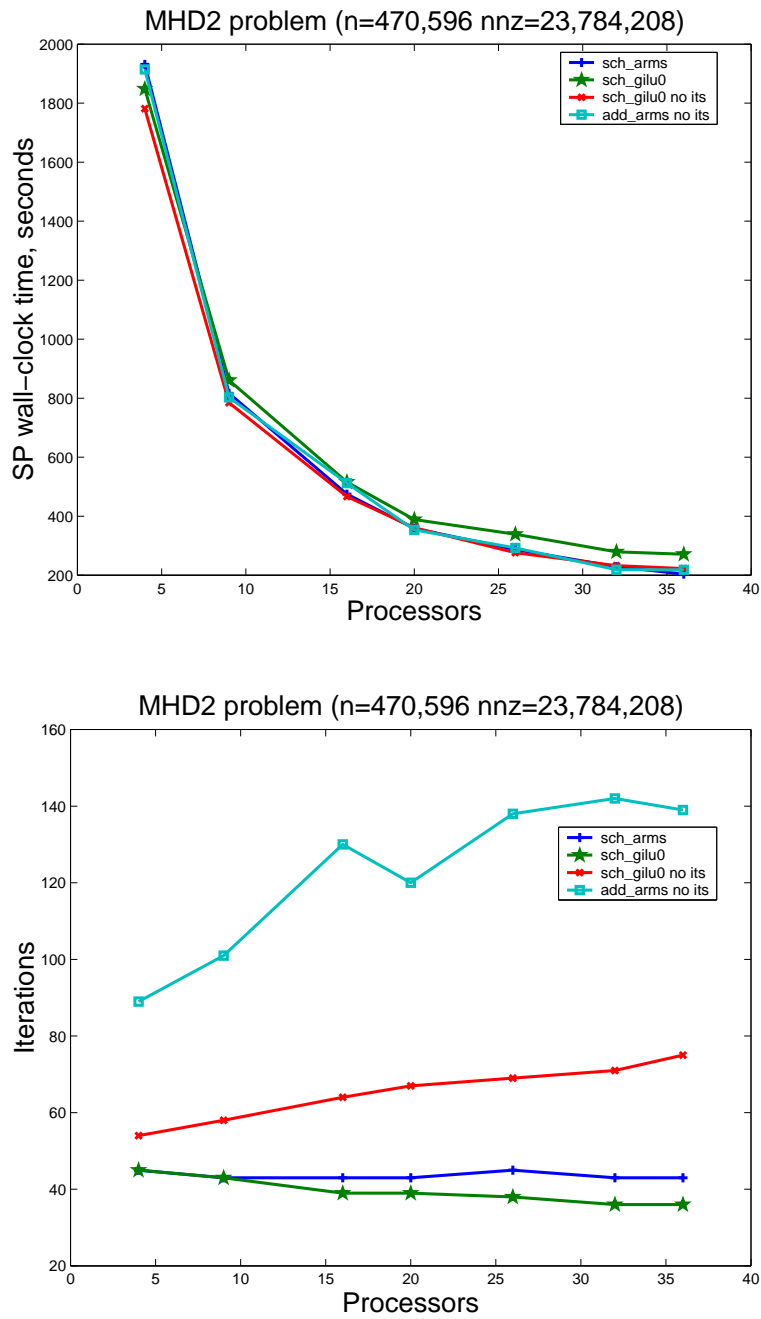


Figure 9. Solution times (top) and outer iterations (bottom) for the (fixed-size) MHD2 problem.

Table IV. Preconditioners in the *hypr* package for the MHD2 problem

Proc.	PILU(1)		PILUT(60,10 <sup>-4</sup> )		PILUT(80,10 <sup>-5</sup> )	
	Iter.	Sol., s	Iter.	Sol., s	Iter.	Sol., s
4	81	1,573.63	769	1,397.71	722	1,570.74
9	103	1,461.50	576	705.45	490	1,226.09

package currently has the most exhaustive selection of parallel preconditioners. In particular, it includes the following options for general sparse linear systems: algebraic multigrid (Boomer-AMG [14]), approximate inverse techniques (ParaSails [33]), a distributed-memory version of parallel ILUT (PILUT [34]), and a parallel level-based ILU(k) (PILU(k) [30]). All of these preconditioners have been tested on four and nine processors, a processor subset used for pARMS, to solve the MHD2 problem, the easier of the two magnetohydrodynamic flow problems. We attempted to construct the preconditioners such that they yield a number of nonzero elements that is comparable with that of the preconditioners tested in the pARMS framework, which is about 2.5 times the number of nonzeros in the original matrix. Note that this actually requires a significant amount of trial and error.

We show the results for the two instances of the PILUT preconditioner. In the first, the fill-in/drop-tolerance parameter pair is taken as 60/10<sup>-4</sup> and in the second, as 80/10<sup>-5</sup>. For parallel ILU(k), the fill level of one gives approximately the desired number of nonzeros. The GMRES restart value and convergence tolerance were set as in the pARMS experiments with this matrix. In ParaSails and Boomer-AMG, the default values for all the preconditioner-specific parameters were taken (see [13] for their values). Table IV shows the obtained results for PILUT and PILU(1). Boomer-AMG did not converge in 1,000 iterations and ParaSails exceeded the three-hour limit imposed on the problem solution time.

The columns *Iter* and *Sol* display the number of iterations to converge and the solution time, respectively. Note that, for PILU(1), the number of iterations is comparable with `add_arms no its` with a better timing on the four processors and almost double the timing on the nine processors. The PILUT(60,10<sup>-4</sup>) factorization has resulted in a much sparser preconditioner than that of preconditioners in pARMS, which appears to be sufficient for convergence but with very large iteration numbers. The high degree of sparsity of the preconditioner results in much less costly preconditioning operations. The more accurate factorization PILUT(80,10<sup>-5</sup>) leads to a reduction in the number of iterations. However, in spite of this reduction, the solution time increases substantially when scaled by the number of processors, when going from four to nine processors. The success of PILUT with a very sparse preconditioner is not too surprising for this is rather easy problem. However, the unexpected big increase of the scaled time for the more realistic case involving a more accurate factorization, raises the question of scalability for this preconditioner. Only a more exhaustive test, involving more problems, and more processors would enable us to reach any conclusion in this regard. Such exhaustive comparisons are beyond the scope of this paper.

## 7. Conclusion

The Algebraic Recursive Multilevel Solvers constitute a viable framework for developing a large array of preconditioners for parallel environments. This framework offers many possibilities ranging from simple Additive Schwarz with or without overlapping, to more complex Schur complement techniques. These techniques are developed for the general sparse irregularly structured linear systems. Experiments confirm what is already known from theory, namely that scalability as well as robustness may suffer when simple Additive Schwarz procedures are used. For harder and larger problems, it is essential to use a truly multilevel method, and this again is already known from general theory. It was also shown that inner iterations at the Schur complement level are essential in this case. While direct solvers are still very competitive for 2-D problems, they are likely to be far inferior, or even to fail, for reasonable size 3-D cases, as was shown in Section 5.3.

Not yet explored in this paper are methods which extend the ARMS recursive reduction into the interdomain interface variables. It remains to be seen whether this further step, which is much more complex to implement, is mandated for much larger problems.

## ACKNOWLEDGEMENTS

We would like to thank Azzeddine Soulaïmani and Ridha Touihri from the “Ecole de Technologie Supérieure, Université du Québec”, for their generous help in supplying us with the realistic test problems of Section 5.5.

## REFERENCES

1. Saad Y. ILUM: a multi-elimination ILU preconditioner for general sparse matrices. *SIAM Journal on Scientific Computing* 1996; **17**(4):830–847.
2. Botta EFF, van der Ploeg A, Wubs FW. A fast linear-system solver for large unstructured problems on a shared-memory computer. In *Proceedings of the Conference on Algebraic Multilevel Methods with Applications*, Axelsson O, Polman B (eds). 1996; 105–116.
3. Botta EFF, Wubs FW. MRILU: it's the preconditioning that counts. *Technical Report W-9703*, Department of Mathematics, University of Groningen, The Netherlands, 1997.
4. Saad Y, Zhang J. BILUM: Block versions of multi-elimination and multi-level ILU preconditioner for general sparse linear systems. *SIAM Journal on Scientific Computing* 1999; **20**:2103–2121.
5. Saad Y, Zhang J. BILUTM: A domain-based multi-level block ILUT preconditioner for general sparse matrices. *SIAM Journal on Matrix Analysis and Applications* 2000; **21**:279–299.
6. Saad Y, Suchomel B. ARMS: an algebraic recursive multilevel solver for general sparse linear systems. *Numerical Linear Algebra with Applications* 2002; **9**:359–378.
7. Saad Y, Sosonkina M, Zhang J. Domain decomposition and multi-level type techniques for general sparse linear systems. In *Domain Decomposition Methods 10*, Mandel J, Farhat C, Cai XC (eds). American Mathematical Society: Providence, 1998; 174–190.
8. Saad Y. *Iterative Methods for Sparse Linear Systems*. PWS publishing: New York, 1996.
9. Bank R, Wagner C. Multilevel ILU decomposition. *Numerische Mathematik* 1999; **82**:543–576.
10. Bank R, Smith RK. An algebraic multilevel multigraph algorithm. *SIAM Journal on Scientific Computing* 2002; **23**:1572–1592.
11. Bröker O, Grote MJ, Mayer C, Reusken A. Robust parallel smoothing for multigrid via sparse approximate inverses. *SIAM Journal on Scientific Computing* 2002; **23**:1396–1417.
12. Notay Y. Using approximate inverses in algebraic multilevel methods. *Numerische Mathematik* 1998; **80**:397–417.
13. Chow E, Cleary A, Falgout R. HYPRE User's manual, version 1.6.0. *Technical Report UCRL-MA-137155*, Lawrence Livermore National Laboratory, Livermore, CA, 1998.



14. Henson VE, Yang UM. BoomerAMG: a parallel algebraic multigrid solver and preconditioner. *Applied Numerical Mathematics* 2002; **41**:155–177.
15. Tuminaro R, Tong C. ML 2.0 smoothed aggregation user's guide. *Technical Report SAND2001-8028*, Sandia National Laboratories, Albuquerque, NM, Dec. 2000.
16. Saad Y, Sasonkina M. Distributed Schur complement techniques for general sparse linear systems. *SIAM Journal on Scientific Computing* 1999; **21**(4):1337–1356.
17. Botta EFF, van der Ploeg A, Wubs FW. Nested grids ILU-decomposition (NGILU). *Journal of Computational and Applied Mathematics* 1996; **66**:515–526.
18. Leuze R. Independent set orderings for parallel matrix factorizations by Gaussian elimination. *Parallel Computing* 1989; **10**:177–191.
19. Wagner C. *Introduction to Algebraic Multigrid - Course Notes of an Algebraic Multigrid Course at the University of Heidelberg in the Wintersemester, 1998/99*.
20. Saad Y, Sasonkina M. Enhanced parallel multicolor preconditioning techniques for linear systems. In *9th SIAM Conference on Parallel Processing for Scientific Computing*. SIAM: Philadelphia, 1999.
21. Bjørstad PE. Multiplicative and Additive Schwarz Methods: Convergence in the 2 domain case. In *Domain Decomposition Methods*, Chan T, Glowinski R, Périaux J, Widlund OB (eds). SIAM: Philadelphia, 1989; 147–159.
22. Bjørstad PE, Widlund OB. To overlap or not to overlap: A note on a domain decomposition method for elliptic problems. *SIAM Journal on Scientific and Statistical Computing* 1989; **10**(5):1053–1061.
23. Cai XC, Gropp WD, Keyes DE. A comparison of some domain decomposition and ILU preconditioned iterative methods for nonsymmetric elliptic problems. *Numerical Linear Algebra with Applications* 1994; **1**:477–504.
24. Gropp WD, Smith BF. Experiences with domain decomposition in three dimensions: Overlapping Schwarz methods. In *Proceedings of the Sixth International Symposium on Domain Decomposition Methods*, Quarteroni A, Périaux J, Kuznetsov YA, Widlund OB (eds). American Mathematical Society, 1993; 323–333.
25. Benzi M, Frommer A, Nabben R, Szyld D. Algebraic theory of overlapping Schwarz methods. *Numerische Mathematik* 2001; **89**:605–639.
26. Cai XC, Sarkis M. A restricted additive Schwarz preconditioner for general sparse linear systems. *SIAM Journal on Scientific Computing* 1999; **21**:792–797.
27. Little L, Li Z, Choi HG, Saad Y. Particle partitioning strategies for the parallel computation of solid-liquid flows. *Computers and Mathematics with Applications* 2002; **43**:1591–1616.
28. Kuznetsov S, Lo GC, Saad Y. Parallel solution of general sparse linear systems using PPARSLIB. In *Domain Decomposition XI*, Lai CH, Bjørstad P, Cross M, Widlund OB (eds). Domain Decomposition Press: Bergen, 1999; 455–465.
29. Soulaïmani A, Rebaine A, Saad Y. Parallelization of the edge-based stabilized finite element method. In *Parallel Computational Fluid Dynamics, Teraflops, Optimization, and novel formulations*. North Holland, 2000; 397–406.
30. Hysom D, Pothen A. A scalable parallel algorithm for incomplete factor preconditioning. *SIAM Journal on Scientific Computing* 2001; **22**(6):2194–2215.
31. Joshi M, Karypis G, Kumar V, Gupta A, Gustavson F. PSPASES: Scalable parallel direct solver library for sparse symmetric positive definite linear systems – user's manual (version 1.0.3). *Technical Report TR-97-059*, Department of Computer Science, University of Minnesota, Minneapolis, MN, 1997.
32. Soulaïmani A, Salah NB, Saad Y. Enhanced GMRES acceleration techniques for some CFD problems. *International Journal of Computational Fluid Dynamics* 2002; **16**(1):1–20.
33. Chow E. *A priori* sparsity patterns for parallel sparse approximate inverse preconditioners. *SIAM Journal on Scientific Computing* 2000; **21**:1804–1822.
34. Karypis G, Kumar V. Parallel threshold-based ilu factorization. *Technical Report TR-96-061*, Department of Computer Science, University of Minnesota, Minneapolis, MN, 1996.