



Parallel sparse linear solvers and applications in CFD

Jocelyne Erhel

Joint work with Désiré Nuentza Wakam (GMRES)
and Baptiste Poirriez (PCG)

SAGE team, Inria Rennes, France



journée Calcul Intensif Distribué dans l'Industrie,
Université Paris 13, 22 janvier 2014



Solver interface

- Interface to direct and iterative solvers: MUMPS, SuperLU_Dist, Hype, Petsc, pArms, etc
- SLSI [Nuentza Wakam et al 2010] available on demand
- System solver in H2OLab platform [Erhel et al 2009]
- Application to CFD problems

GMRES(m): a Krylov method for general matrices

- combining Domain Decomposition and deflation

PCG: a Krylov method for SPD matrices

- combining Domain Decomposition and deflation



$$Ax = b, \quad A \in \mathbb{R}^{n \times n} \quad x, b \in \mathbb{R}^n \quad B = AM^{-1}$$

GMRES(m): a Krylov subspace method

- [Saad and Schultz 1986, Meurant's book 1999, Saad's book 2003, Simoncini and Szyld 2007, Erhel 2011, ...]
- Fix x_0 , then $r_0 = b - Ax_0$
- $\mathcal{K}_m(B, r_0) = \text{span}\{r_0, Br_0, \dots, B^{m-1}r_0\}$
- Find $x_m \in x_0 + \mathcal{K}_m(B, r_0)$ such that $\|r_m\|_2 = \|b - Bx_m\|_2 = \min_{u \in x_0 + \mathcal{K}_m(B, r_0)} \|b - Bu\|_2$

Building blocks of GMRES

- Initial step: choose x_0 , compute r_0
- First step: generate an orthonormal basis $V_{m+1} = [v_0, \dots, v_m]$ of $\mathcal{K}_{m+1}(B, r_0)$ such that
$$v_0 = r_0/\beta, \quad \beta = \|r_0\|_2, \quad BV_m = V_{m+1}\bar{H}_m$$
- Second step: approximate solution $x_m = x_0 + M^{-1}V_my_m$
$$\Rightarrow r_m = r_0 - BV_my_m = V_{m+1}(\beta e_1 - \bar{H}_my_m)$$
$$\Rightarrow y_m = \min_{y \in \mathbb{R}^m} \|\beta e_1 - \bar{H}_my\|_2$$



Arnoldi process

```

1:  $v_0 = r_0 / \|r_0\|_2$ 
2: for  $k = 0, \dots$  do
3:    $p = Bv_k$ 
4:   for  $i = 1 : k$  do
5:      $h_{ik} = v_i^T p$ 
6:      $p = p - h_{ik} v_i$ 
7:   end for
8:    $h_{k+1,k} = \|p\|_2$ 
9:    $v_{k+1} = p / h_{k+1,k}$ 
10: end for
    
```



$$BV_m = V_{m+1} \tilde{H}_m$$

Granularity issues in parallel algorithms

⇒ Communication-avoiding strategies

- Generate the basis vectors [Reichel 1990, Bai et al 1994]
- Orthogonalize the basis [De Sturler 1994, Erhel 1995, Sidje 1997]
- Improve the strategy [Hoemmen 2010, Demmel et al 2011]

Preconditioning issues

⇒ use multilevel methods to deal with large systems

- Schwarz preconditioning [Atenekeng Kahou et al 2007, Dufaud+Tromeur-Dervout 2010, Giraud+Haidar 2009, Smith et al's book 1996,...]
- Filtering and Schur complement [Li et al 2003, Grigori et al 2011]
- Multilevel parallelism [Nuentsu Wakam et al 2011, Giraud et al 2010, ...]

Complexity and stagnation issues with restarted GMRES(m)

⇒ Use deflation to recover possible loss of information

- Deflation by preconditioning [Erhel et al 1996, Burrage et al 1998, Baglama et al 1998, ...]
- Deflation by augmented basis [Morgan 1995, Morgan 2002,...]

Work in the team SAGE

Combine 'communication-avoiding' GMRES ... and Deflation ... and domain decomposition preconditioners
[Nuentsu Wakam 2011, Nuentsu Wakam+Erhel+Gropp 2013, Nuentsu Wakam+Pacull 2013, Nuentsu Wakam+Erhel 2014]



Arnoldi process

```

1:  $v_0 = r_0 / \|r_0\|_2$ 
2: for  $k = 0, \dots$  do
3:    $p = Bv_k$ 
4:   for  $i = 1 : k$  do
5:      $h_{ik} = v_i^T p$ 
6:      $p = p - h_{ik} v_i$ 
7:   end for
8:    $h_{k+1,k} = \|p\|_2$ 
9:    $v_{k+1} = p / h_{k+1,k}$ 
10: end for
    
```



$$BV_m = V_{m+1} \tilde{H}_m$$

Granularity issues in parallel algorithms

⇒ Communication-avoiding strategies

- Generate the basis vectors [Reichel 1990, Bai et al 1994]
- Orthogonalize the basis [De Sturler 1994, Erhel 1995, Sidje 1997]
- Improve the strategy [Hoemmen 2010, Demmel et al 2011]

Preconditioning issues

⇒ use multilevel methods to deal with large systems

- Schwarz preconditioning [Atenekeng Kahou et al 2007, Dufaud+Tromeur-Dervout 2010, Giraud+Haidar 2009, Smith et al's book 1996,...]
- Filtering and Schur complement [Li et al 2003, Grigori et al 2011]
- Multilevel parallelism [Nuentsa Wakam et al 2011, Giraud et al 2010, ...]

Complexity and stagnation issues with restarted GMRES(m)

⇒ Use deflation to recover possible loss of information

- Deflation by preconditioning [Erhel et al 1996, Burrage et al 1998, Baglama et al 1998, ...]
- Deflation by augmented basis [Morgan 1995, Morgan 2002,...]

Work in the team SAGE

Combine 'communication-avoiding' GMRES ... and Deflation ... and domain decomposition preconditioners [Nuentsa Wakam 2011, Nuentsa Wakam+Erhel+Gropp 2013, Nuentsa Wakam+Pacull 2013, Nuentsa Wakam+Erhel 2014]



Arnoldi process

```

1:  $v_0 = r_0 / \|r_0\|_2$ 
2: for  $k = 0, \dots$  do
3:    $p = Bv_k$ 
4:   for  $i = 1 : k$  do
5:      $h_{ik} = v_i^T p$ 
6:      $p = p - h_{ik} v_i$ 
7:   end for
8:    $h_{k+1,k} = \|p\|_2$ 
9:    $v_{k+1} = p / h_{k+1,k}$ 
10: end for
    
```



$$BV_m = V_{m+1} \tilde{H}_m$$

Granularity issues in parallel algorithms

⇒ Communication-avoiding strategies

- Generate the basis vectors [Reichel 1990, Bai et al 1994]
- Orthogonalize the basis [De Sturler 1994, Erhel 1995, Sidje 1997]
- Improve the strategy [Hoemmen 2010, Demmel et al 2011]

Preconditioning issues

⇒ use multilevel methods to deal with large systems

- Schwarz preconditioning [Atenekeng Kahou et al 2007, Dufaud+Tromeur-Dervout 2010, Giraud+Haidar 2009, Smith et al's book 1996,...]
- Filtering and Schur complement [Li et al 2003, Grigori et al 2011]
- Multilevel parallelism [Nuentsa Wakam et al 2011, Giraud et al 2010, ...]

Complexity and stagnation issues with restarted GMRES(m)

⇒ Use deflation to recover possible loss of information

- Deflation by preconditioning [Erhel et al 1996, Burrage et al 1998, Baglama et al 1998, ...]
- Deflation by augmented basis [Morgan 1995, Morgan 2002,...]

Work in the team SAGE

Combine 'communication-avoiding' GMRES ... and Deflation ... and domain decomposition preconditioners [Nuentsa Wakam 2011, Nuentsa Wakam+Erhel+Gropp 2013, Nuentsa Wakam+Pacull 2013, Nuentsa Wakam+Erhel 2014]



Arnoldi process

```

1:  $v_0 = r_0 / \|r_0\|_2$ 
2: for  $k = 0, \dots$  do
3:    $p = Bv_k$ 
4:   for  $i = 1 : k$  do
5:      $h_{ik} = v_i^T p$ 
6:      $p = p - h_{ik} v_i$ 
7:   end for
8:    $h_{k+1,k} = \|p\|_2$ 
9:    $v_{k+1} = p / h_{k+1,k}$ 
10: end for
    
```



$$BV_m = V_{m+1} \tilde{H}_m$$

Granularity issues in parallel algorithms

⇒ Communication-avoiding strategies

- Generate the basis vectors [Reichel 1990, Bai et al 1994]
- Orthogonalize the basis [De Sturler 1994, Erhel 1995, Sidje 1997]
- Improve the strategy [Hoemmen 2010, Demmel et al 2011]

Preconditioning issues

⇒ use multilevel methods to deal with large systems

- Schwarz preconditioning [Atenekeng Kahou et al 2007, Dufaud+Tromeur-Dervout 2010, Giraud+Haidar 2009, Smith et al's book 1996,...]
- Filtering and Schur complement [Li et al 2003, Grigori et al 2011]
- Multilevel parallelism [Nuentsu Wakam et al 2011, Giraud et al 2010, ...]

Complexity and stagnation issues with restarted GMRES(m)

⇒ Use deflation to recover possible loss of information

- Deflation by preconditioning [Erhel et al 1996, Burrage et al 1998, Baglama et al 1998, ...]
- Deflation by augmented basis [Morgan 1995, Morgan 2002,...]

Work in the team SAGE

Combine 'communication-avoiding' GMRES ... and Deflation ... and domain decomposition preconditioners [Nuentsu Wakam 2011, Nuentsu Wakam+Erhel+Gropp 2013, Nuentsu Wakam+Pacull 2013, Nuentsu Wakam+Erhel 2014]



Arnoldi process

```

1:  $v_0 = r_0 / \|r_0\|_2$ 
2: for  $k = 0, \dots$  do
3:    $p = Bv_k$ 
4:   for  $i = 1 : k$  do
5:      $h_{ik} = v_i^T p$ 
6:      $p = p - h_{ik} v_i$ 
7:   end for
8:    $h_{k+1,k} = \|p\|_2$ 
9:    $v_{k+1} = p / h_{k+1,k}$ 
10: end for
    
```



$$BV_m = V_{m+1} \tilde{H}_m$$

Granularity issues in parallel algorithms

⇒ Communication-avoiding strategies

- Generate the basis vectors [Reichel 1990, Bai et al 1994]
- Orthogonalize the basis [De Sturler 1994, Erhel 1995, Sidje 1997]
- Improve the strategy [Hoemmen 2010, Demmel et al 2011]

Preconditioning issues

⇒ use multilevel methods to deal with large systems

- Schwarz preconditioning [Atenekeng Kahou et al 2007, Dufaud+Tromeur-Dervout 2010, Giraud+Haidar 2009, Smith et al's book 1996,...]
- Filtering and Schur complement [Li et al 2003, Grigori et al 2011]
- Multilevel parallelism [Nuentsa Wakam et al 2011, Giraud et al 2010, ...]

Complexity and stagnation issues with restarted GMRES(m)

⇒ Use deflation to recover possible loss of information

- Deflation by preconditioning [Erhel et al 1996, Burrage et al 1998, Baglama et al 1998, ...]
- Deflation by augmented basis [Morgan 1995, Morgan 2002,...]

Work in the team SAGE

Combine 'communication-avoiding' GMRES ... and Deflation ... and domain decomposition preconditioners
[Nuentsa Wakam 2011, Nuentsa Wakam+Erhel+Gropp 2013, Nuentsa Wakam+Pacull 2013, Nuentsa Wakam+Erhel 2014]



building blocks

- Initial step: run one cycle of GMRES(m) and compute shifts for the Newton basis
- First step: build a basis $K_{m+1} = [k_0, k_1, \dots, k_m]$ of the Krylov subspace $\mathcal{K}_{m+1}(B, r_0)$ such that

$$BK_m = K_{m+1} \bar{T}_m$$

- Second step: compute an orthonormal basis of $\mathcal{K}_{m+1}(B, r_0)$
Compute the QR factorization $K_{m+1} = V_{m+1} R_{m+1}$
RODDEC [Sidje 1997, Erhel 1995] or TSQR [Demmel et al 2011]

$$\Rightarrow BK_m = V_{m+1} R_{m+1} \bar{T}_m \Rightarrow BV_m = V_{m+1} \underbrace{R_{m+1} \bar{T}_m R_m^{-1}}_{\bar{H}_m}$$

- Third step: approximate solution $x_m = x_0 + M^{-1} V_m y_m$

$$\Rightarrow r_m = r_0 - BK_m y_m = V_{m+1} (\beta e_1 - \bar{H}_m y) \quad \text{with } \beta = \|r_0\|_2$$

$$\Rightarrow y_m = \min_{y \in \mathbb{R}^m} \|\beta e_1 - \bar{H}_m y\|_2$$



building blocks

- Initial step: run one cycle of GMRES(m) and compute shifts for the Newton basis
- First step: build a basis $K_{m+1} = [k_0, k_1, \dots, k_m]$ of the Krylov subspace $\mathcal{K}_{m+1}(B, r_0)$ such that

$$BK_m = K_{m+1} \bar{T}_m$$

- Second step: compute an orthonormal basis of $\mathcal{K}_{m+1}(B, r_0)$
Compute the QR factorization $K_{m+1} = V_{m+1} R_{m+1}$
RODDEC [Sidje 1997, Erhel 1995] or TSQR [Demmel et al 2011]

$$\Rightarrow BK_m = V_{m+1} R_{m+1} \bar{T}_m \Rightarrow BV_m = V_{m+1} \underbrace{R_{m+1} \bar{T}_m R_m^{-1}}_{\bar{H}_m}$$

- Third step: approximate solution $x_m = x_0 + M^{-1} V_m y_m$

$$\Rightarrow r_m = r_0 - BK_m y_m = V_{m+1} (\beta e_1 - \bar{H}_m y_m) \quad \text{with } \beta = \|r_0\|_2$$

$$\Rightarrow y_m = \min_{y \in \mathbb{R}^m} \|\beta e_1 - \bar{H}_m y\|_2$$



building blocks

- Initial step: run one cycle of GMRES(m) and compute shifts for the Newton basis
- First step: build a basis $K_{m+1} = [k_0, k_1, \dots, k_m]$ of the Krylov subspace $\mathcal{K}_{m+1}(B, r_0)$ such that

$$BK_m = K_{m+1} \bar{T}_m$$

- Second step: compute an orthonormal basis of $\mathcal{K}_{m+1}(B, r_0)$
Compute the QR factorization $K_{m+1} = V_{m+1} R_{m+1}$
RODDEC [Sidje 1997, Erhel 1995] or TSQR [Demmel et al 2011]

$$\Rightarrow BK_m = V_{m+1} R_{m+1} \bar{T}_m \Rightarrow BV_m = V_{m+1} \underbrace{R_{m+1} \bar{T}_m R_m^{-1}}_{\bar{H}_m}$$

- Third step: approximate solution $x_m = x_0 + M^{-1} V_m y_m$

$$\Rightarrow r_m = r_0 - BK_m y_m = V_{m+1} (\beta e_1 - \bar{H}_m y_m) \quad \text{with } \beta = \|r_0\|_2$$

$$\Rightarrow y_m = \min_{y \in \mathbb{R}^m} \|\beta e_1 - \bar{H}_m y\|_2$$



building blocks

- Initial step: run one cycle of GMRES(m) and compute shifts for the Newton basis
- First step: build a basis $K_{m+1} = [k_0, k_1, \dots, k_m]$ of the Krylov subspace $\mathcal{K}_{m+1}(B, r_0)$ such that

$$BK_m = K_{m+1} \bar{T}_m$$

- Second step: compute an orthonormal basis of $\mathcal{K}_{m+1}(B, r_0)$
Compute the QR factorization $K_{m+1} = V_{m+1} R_{m+1}$
RODDEC [Sidje 1997, Erhel 1995] or TSQR [Demmel et al 2011]

$$\Rightarrow BK_m = V_{m+1} R_{m+1} \bar{T}_m \Rightarrow BV_m = V_{m+1} \underbrace{R_{m+1} \bar{T}_m R_m^{-1}}_{\bar{H}_m}$$

- Third step: approximate solution $x_m = x_0 + M^{-1} V_m y_m$

$$\Rightarrow r_m = r_0 - BK_m y_m = V_{m+1} (\beta e_1 - \bar{H}_m y) \quad \text{with } \beta = \|r_0\|_2$$

$$\Rightarrow y_m = \min_{y \in \mathbb{R}^m} \|\beta e_1 - \bar{H}_m y\|_2$$



Main steps with Domain Decomposition preconditioning

- Partition the weighted graph of the matrix in parallel with PARMETIS.
- Redistribute the matrix and right-hand-side according to the PARMETIS partitioning.
- Perform a parallel iterative row and column scaling on the matrix and the right-hand side vector [Amestoy et al, 2008].
- Define the overlap between the submatrices for the additive Schwarz preconditioner.

$$M_{RAS}^{-1} = \sum_{k=1}^D (R_k^0)^T (A_k^\delta)^{-1} R_k^\delta$$

- Setup the submatrices (ILU or LU factorization).
- Solve iteratively the preconditioned system using GMRES.



Restarted GMRES(m)

- $x_m = x_0 + M^{-1}V_m y_m$ where y_m minimizes $\|r_m\|_2$
- The convergence rate depends on the spectral distribution in B
- Smallest eigenvalues slow down the convergence
- Deflation occurs when the Krylov subspace is large enough
- With restarting : loss of spectral information, risk of stalling

Accelerating the restarted GMRES [Simoncini and Szyld, 2007]

- Approximate the smallest eigenvalues and the associated invariant subspace U_r
- Explicit deflation technique [Erhel et al 1996; Burrage et al 1998; Moriya et al 2000]:

$$B\bar{M}^{-1}\bar{x} = b$$

with $\bar{M}^{-1} = (I_n + U_r(|\lambda_n|T^{-1} - I_r)U_r^T$ and $T = U_r^T B U_r$,

- Augmented techniques [Morgan 2000, 2002, Giraud et al, 2010]:

$$x_m \in x_0 + \text{span}\{U_r\} + \mathcal{K}_m(B, r_0)$$



Restarted GMRES(m)

- $x_m = x_0 + M^{-1}V_m y_m$ where y_m minimizes $\|r_m\|_2$
- The convergence rate depends on the spectral distribution in B
- Smallest eigenvalues slow down the convergence
- Deflation occurs when the Krylov subspace is large enough
- With restarting : loss of spectral information, risk of stalling

Accelerating the restarted GMRES [Simoncini and Szyld, 2007]

- Approximate the smallest eigenvalues and the associated invariant subspace U_r
- Explicit deflation technique [Erhel et al 1996; Burrage et al 1998; Moriya et al 2000]:

$$B\bar{M}^{-1}\bar{x} = b$$

with $\bar{M}^{-1} = (I_n + U_r(|\lambda_n|T^{-1} - I_r)U_r^T$ and $T = U_r^T B U_r$

- Augmented techniques [Morgan 2000, 2002, Giraud et al, 2010]:

$$x_m \in x_0 + \text{span}\{U_r\} + \mathcal{K}_m(B, r_0)$$



DGMRES(m, r)

- Perform one cycle of restarted GMRES(m) and compute shifts for the Newton basis
- compute U_r , a basis of a coarse subspace
- Build $\tilde{M}_r^{-1} \equiv I_n + U_r(|\lambda_n|T^{-1} - I_r)U_r^T$, $T = U_r^T B U_r$
- Apply GMRES(m) to $B\tilde{M}_r^{-1}$
- At each restart, update r and the basis U_r

Adaptive DGMRES(m, r)

- Detect a potential slow convergence [Sosonkina et al 1998]
- Switch to DGMRES(m, r) only if necessary [Nüentsa Wakam et al 2013]



DGMRES(m, r)

- Perform one cycle of restarted GMRES(m) and compute shifts for the Newton basis
- compute U_r , a basis of a coarse subspace
- Build $\tilde{M}_r^{-1} \equiv I_n + U_r(|\lambda_n|T^{-1} - I_r)U_r^T$, $T = U_r^T B U_r$
- Apply GMRES(m) to $B\tilde{M}_r^{-1}$
- At each restart, update r and the basis U_r

Adaptive DGMRES(m, r)

- Detect a potential slow convergence [Sosonkina et al 1998]
- Switch to DGMRES(m, r) only if necessary [Nuentza Wakam et al 2013]



DGMRES(m, r)

- Perform one cycle of restarted GMRES(m) and compute shifts for the Newton basis
- compute U_r , a basis of a coarse subspace
- Build $\tilde{M}_r^{-1} \equiv I_n + U_r(|\lambda_n|T^{-1} - I_r)U_r^T$, $T = U_r^T B U_r$
- Apply GMRES(m) to $B\tilde{M}_r^{-1}$
- At each restart, update r and the basis U_r

Adaptive DGMRES(m, r)

- Detect a potential slow convergence [Sosonkina et al 1998]
- Switch to DGMRES(m, r) only if necessary [Nuentza Wakam et al 2013]

Implementation of Deflated GMRES in PETSc



SAGE-
solvers

JE

GMRES

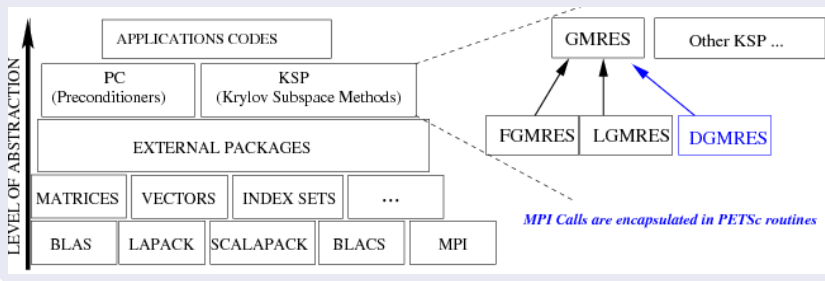
issues

DGMRES

AGMRES

PCG

New KSP type : DGMRES



Usage in Petsc

- Available in PETSc
- Use DGMRES just as any other KSP with the following options




FLUOREM matrices


- in MatrixMarket collection
- large, sparse, nonsymmetric matrices
- linearization of Navier-Stokes: symmetric profile with structured blocks
- Schwarz preconditioning combined with GMRES or DGMRES

[Nuentsa Wakam+Erhel+Gropp 2013; Nuentsa Wakam+Pacull 2013]

RM07R $n= 381,689$; $nnz=37,464,962$



HV15R $n= 2,017,169$; $nnz=283,073,458$



Combining DGMRES with domain decomposition



SAGE-
solvers

JE

GMRES


issues

DGMRES


AGMRES

PCG

RM07R



HV15R



CPU time on parallel computers

RM07R, $n = 381,689$, $nz = 37,464,962$							
D	FULL-GMRES		GMRES(48)		DGMRES(32,2,15)		
	ITS	Time (s)	ITS	Time (s)	ITS	Time (s)	r
16	92	214	169	297	115	250	9
32	117	103	355	260	160	139	11
64	149	66	860	166	206	53	12

HV15R, $n = 2,017,169$, $nz = 283,073,458$							
D	FULL-GMRES		GMRES(150)		DGMRES(96,10,70)		
	ITS	Time (s)	ITS	Time (s)	ITS	Time (s)	r
32	389	1536	(8.4E-08)	3,316	510	2002	50
64	494	1024	(6.8E-01)	-	635	1500	61



Building blocks

- Initial step: run one cycle of GMRES(m) and compute shifts for the Newton basis
Compute $U_r = [u_0, u_1, \dots, u_{r-1}]$ a basis of a coarse subspace
- First step: build a basis $K_{m+1} = [k_0, k_1, \dots, k_m]$ of the Krylov subspace $\mathcal{K}_{m+1}(B, r_0)$ such that

$$BK_m = K_{m+1} \bar{T}_m$$

Define the augmented subspace $\mathcal{C}_s = \mathcal{K}_m(B, r_0) + \text{span}\{U_r\}$ with $s = m + r$ with the basis

$$\begin{bmatrix} K_m & U_r \end{bmatrix}$$

compute

$$BU_r = \hat{K}_r D_r$$

Define the augmented subspace $\hat{\mathcal{C}}_{s+1} = \mathcal{K}_{m+1}(B, r_0) + \text{span}\{BU_r\}$ with the basis

$$\begin{bmatrix} K_{m+1} & \hat{K}_r \end{bmatrix}$$



Building blocks

- Initial step: run one cycle of GMRES(m) and compute shifts for the Newton basis
Compute $U_r = [u_0, u_1, \dots, u_{r-1}]$ a basis of a coarse subspace
- First step: build a basis $K_{m+1} = [k_0, k_1, \dots, k_m]$ of the Krylov subspace $\mathcal{K}_{m+1}(B, r_0)$ such that

$$BK_m = K_{m+1} \bar{T}_m$$

Define the augmented subspace $\mathcal{C}_s = \mathcal{K}_m(B, r_0) + \text{span}\{U_r\}$ with $s = m + r$ with the basis

$$\begin{bmatrix} K_m & U_r \end{bmatrix}$$

compute

$$BU_r = \hat{K}_r D_r$$

Define the augmented subspace $\hat{\mathcal{C}}_{s+1} = \mathcal{K}_{m+1}(B, r_0) + \text{span}\{BU_r\}$ with the basis

$$\begin{bmatrix} K_{m+1} & \hat{K}_r \end{bmatrix}$$



Building blocks

- Second step: Compute an orthonormal basis of \hat{C}_{s+1}

QR factorize the augmented basis $\begin{bmatrix} K_{m+1} & \hat{K}_r \end{bmatrix} = V_{s+1} R_{s+1}$

$$\Rightarrow BK_m = V_{m+1} R_{m+1} \bar{T}_m \Rightarrow BV_m = V_{m+1} R_{m+1} \bar{T}_m R_m^{-1}$$

$$\Rightarrow BU_r = (V_{m+1} R_{m+1,r} + V_r R_r) D_r$$

Define the basis $W_s = \begin{bmatrix} V_m & U_r \end{bmatrix}$

$$\Rightarrow BW_s = V_{s+1} \bar{H}_s$$

- Third step: $x_s = x_0 + M^{-1} W_s y_s$

$$\Rightarrow r_s = r_0 - BW_s y_s = V_{s+1} (\beta e_1 - \bar{H}_s y_s) \quad \text{and } \beta = \|r_0\|_2$$

$$y_s = \min_{y \in \mathbb{R}^s} \|\beta e_1 - \bar{H}_s y\|_2$$

- Final step: Adaptively update r and the coarse basis U_r



Building blocks

- Second step: Compute an orthonormal basis of \hat{C}_{s+1}

QR factorize the augmented basis $\begin{bmatrix} K_{m+1} & \hat{K}_r \end{bmatrix} = V_{s+1} R_{s+1}$

$$\Rightarrow BK_m = V_{m+1} R_{m+1} \bar{T}_m \Rightarrow BV_m = V_{m+1} R_{m+1} \bar{T}_m R_m^{-1}$$

$$\Rightarrow BU_r = (V_{m+1} R_{m+1,r} + V_r R_r) D_r$$

Define the basis $W_s = \begin{bmatrix} V_m & U_r \end{bmatrix}$

$$\Rightarrow BW_s = V_{s+1} \bar{H}_s$$

- Third step: $x_s = x_0 + M^{-1} W_s y_s$

$$\Rightarrow r_s = r_0 - BW_s y_s = V_{s+1} (\beta e_1 - \bar{H}_s y_s) \quad \text{and } \beta = \|r_0\|_2$$

$$y_s = \min_{y \in \mathbb{R}^s} \|\beta e_1 - \bar{H}_s y\|_2$$

- Final step: Adaptively update r and the coarse basis U_r



Building blocks

- Second step: Compute an orthonormal basis of \hat{C}_{s+1}

QR factorize the augmented basis $\begin{bmatrix} K_{m+1} & \hat{K}_r \end{bmatrix} = V_{s+1} R_{s+1}$

$$\Rightarrow BK_m = V_{m+1} R_{m+1} \bar{T}_m \Rightarrow BV_m = V_{m+1} R_{m+1} \bar{T}_m R_m^{-1}$$

$$\Rightarrow BU_r = (V_{m+1} R_{m+1,r} + V_r R_r) D_r$$

Define the basis $W_s = \begin{bmatrix} V_m & U_r \end{bmatrix}$

$$\Rightarrow BW_s = V_{s+1} \bar{H}_s$$

- Third step: $x_s = x_0 + M^{-1} W_s y_s$

$$\Rightarrow r_s = r_0 - BW_s y_s = V_{s+1} (\beta e_1 - \bar{H}_s y_s) \quad \text{and } \beta = \|r_0\|_2$$


$$y_s = \min_{y \in \mathbb{R}^s} \|\beta e_1 - \bar{H}_s y\|_2$$

- Final step: Adaptively update r and the coarse basis U_r



New KSP type : AGMRES

figure



Usage in Petsc

- Use AGMRES just as GMRES
- \Rightarrow `KSPSetType(ksp, KSPAGMRES)` or `-ksp_type agmres, -pc_type asm, ...`
- Options : `-ksp_gmres_restart m, -ksp_agmres_eig r,`
- `-ksp_max_its maxits, -ksp_agmres_smv smv -ksp_agmres_bgv bgv, ...`

Experiments with Augmented GMRES



SAGE-
solvers

JE

GMRES


issues

DGMRES


AGMRES

PCG

RM07R: effect of the restarting



RM07R: effect of the number of subdomains



CPU time on parallel computers

RM07R, $n = 381,689$, $nz = 37,464,962$				
D	GMRES(32)		AGMRES(32,r)	
	ITS	Time (s)	ITS	Time (s)
16	254	379.3	169	224.1
32	886	573.4	212	91.41
64	-	-	287	62.39

Parallel CPU Time with AGMRES



SAGE-
solvers

JE

GMRES

issues

DGMRES

AGMRES


PCG

Convection-Diffusion test cases


- 3DCONSKY_121 : size = 1,771,561; nonzeros = 50,178,241
- 3DCOSKY_161 : size= 4,173,281; nonzeros = 118,645,121

[Nuentza Wakam+Erhel 2014]

3DCONSKY_121



3DCONSKY_161





PCG

- A Symmetric Positive Definite (SPD) matrix
- PCG is a Krylov method
- short recurrences and minimization properties
- preconditioning M^{-1}

Coarse grid and balancing

[Nicolaides 1987, Mandel 1993, DD proceedings, Giraud et al.]

- Z basis of a coarse subspace
- restriction of A : nonsingular small matrix $A_c = Z^T A Z$
- projections: $P = I - A Z A_c^{-1} Z^T$ and $P^T = I - Z A_c^{-1} (A Z)^T$
- coarse grid: $Z A_c^{-1} Z^T$
- balancing: $C_b = P^T M^{-1} P + Z A_c^{-1} Z^T$

Coarse grid and augmented CG

[Erhel et al 2000, Saad et al. 2000, Tang et al. 2009, Poirriez 2011, Nataf et al]

- $x_0 = Z A_c^{-1} Z^T b$
- $C_a = P^T M^{-1}$
- C_a is equivalent to C_b (if no loss of orthogonality)



Balancing Neumann Neumann

- PCG applied to a Schur complement
- Neumann-Neumann preconditioning M^{-1}
- Balancing with a coarse grid Z

SIDNUR

[Poirriez 2011, Pichot et al. 2014]

- domain decomposition provided by the user
- coarse grid : signature of subdomains [Frank and Vuik 2001]
- C++ library
- mutual factorization of local Schur complements and local matrices
- management of floating subdomains
- numerical experiments with 3D fracture networks




Flow computations in Discrete Fracture Networks


- random domain generated with MPFRAC software
- SPD sparse matrix
- solving with SIDNUR

[Poirriez 2011]

Number of iterations



CPU time

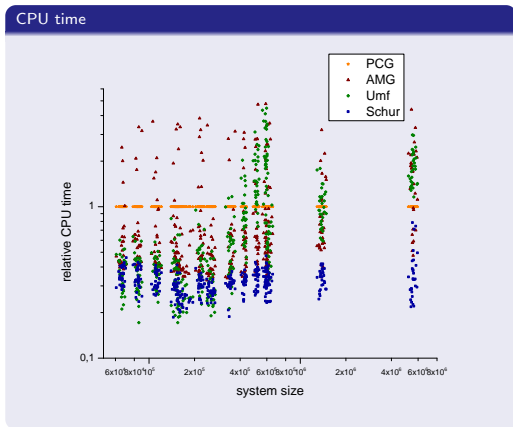


Comparing SIDNUR with other solvers



- UMFPACK: direct solver
- Boomer-AMG: algebraic multigrid
- PCG: Conjugate Gradient preconditioned by Boomer-AMG

[Poirriez 2011]



SAGE-
solvers

JE

GMRES

PCG

SIDNUR



GMRES

- DGMRES KSP module: deflation in GMRES(m) with or without Newton basis
- AGMRES KSP module: augmented Newton basis in GMRES(m)
- Deflation combined with Schwarz domain decomposition preconditioning
- Robustness: reduce the restarting effects and the domain decomposition effects
- Efficiency: increase granularity and scalability
- Numerical experiments with CFD problems: DGMRES and AGMRES faster than GMRES

PCG

- Deflation combined with Schur domain decomposition
- SIDNUR: Balancing Domain Decomposition
- Robustness: reduce the domain decomposition effects
- Efficiency: parallel Schur and Neumann Neumann computations
- Numerical experiments with 3D fracture networks: faster than multigrid and PCG
- Library soon available as free software