Krylov subspace methods: CG

L. Grigori

EPFL and PSI

December 3, 2024

Plan

Plan

Krylov subspace methods

- Symmetric matrices and Lanczos algorithm
- Conjugate gradient method

Enlarged Krylov methods

Krylov subspace methods for solving $Ax = b$

Finds a sequence $x_1, x_2, ..., x_m$ that minimizes some measure of error over the spaces $x_0 + \mathcal{K}_i(A, r_0),\ i = 1, ..., m$, by satisfying two conditions:

1. Subspace condition: $x_m \in x_0 + \mathcal{K}_m(A, r_0)$
2. Petrov-Galerkin condition: $r_m \perp \mathscr{L}_m \iff (r_m)^t y = 0,\ \forall\ y \in \mathscr{L}_m$

where

- x_0 initial iterate, r_0 initial residual, $\mathscr{L}_m$ well-defined subspace of dimension m
- $\mathcal{K}_m(A, r_0) = span\{r_0, Ar_0, A^2 r_0, ..., A^{m-1} r_0\}$ is the Krylov subspace of dimension m.

Two instances of Krylov projection method:

Conjugate gradient [Hestenes, Stieffel, 52], A is SPD, $\mathscr{L}_m = \mathcal{K}_m(A, r_0)$,

- finds x_m by minimizing $\|x - x_m\|_A$ over $x_0 + \mathcal{K}_m(A, r_0)$

GMRES [Saad, Schultz, 86], A is unsymmetric, $\mathscr{L}_m = A\mathcal{K}_m(A, r_0)$,

- finds x_m by minimizing $\|Ax - b\|_2$ over $x_0 + \mathcal{K}_m(A, r_0)$.

Krylov subspace methods for solving $Ax = b$

Finds a sequence $x_1, x_2, ..., x_m$ that minimizes some measure of error over the spaces $x_0 + \mathcal{K}_i(A, r_0),\ i = 1, ..., m$, by satisfying two conditions:

1. Subspace condition: $x_m \in x_0 + \mathcal{K}_m(A, r_0)$
2. Petrov-Galerkin condition: $r_m \perp \mathscr{L}_m \iff (r_m)^t y = 0,\ \forall\ y \in \mathscr{L}_m$

where

- x_0 initial iterate, r_0 initial residual, $\mathscr{L}_m$ well-defined subspace of dimension m
- $\mathcal{K}_m(A, r_0) = span\{r_0, Ar_0, A^2 r_0, ..., A^{m-1} r_0\}$ is the Krylov subspace of dimension m.

Two instances of Krylov projection method:

Conjugate gradient [Hestenes, Stieffel, 52], A is SPD, $\mathscr{L}_m = \mathcal{K}_m(A, r_0)$,

- finds x_m by minimizing $\|x - x_m\|_A$ over $x_0 + \mathcal{K}_m(A, r_0)$

GMRES [Saad, Schultz, 86], A is unsymmetric, $\mathscr{L}_m = A\mathcal{K}_m(A, r_0)$,

- finds x_m by minimizing $\|Ax - b\|_2$ over $x_0 + \mathcal{K}_m(A, r_0)$.

Orthogonal Projection Methods (OPM)

Finds a sequence $x_1, x_2, ..., x_m$ that minimizes some measure of error over the spaces $x_0 + \mathcal{K}_i(A, r_0),\ i = 1, ..., m$, by satisfying two conditions:

1. Subspace condition: $x_m \in x_0 + \mathcal{K}_m(A, r_0)$
2. Sketched Petrov-Galerkin condition:
   $r_m \perp \mathcal{K}_m \iff (r_m)^t y = 0,\ \forall\ y \in \mathcal{K}_m$

where

- x_0 initial iterate, r_0 initial residual

- $\mathcal{K}_m(A, r_0) = span\{r_0, Ar_0, A^2 r_0, ..., A^{m-1} r_0\}$ is the Krylov subspace of dimension m.

Optimality result for Orthogonal Projection Methods

Proposition Assume A is Symmetric Positive Definite. Then $\tilde{x}$ is the result of an (orthogonal) projection method onto $\mathcal{K}$ with the starting vector x_0 if and only if it minimizes the A-norm of the error over $x_0 + \mathcal{K}$, i.e.,

$$E(\tilde{x}) = \min_{x \in x_0 + \mathcal{K}} E(x),$$

where $x^* = A^{-1}b$ and

$$E(x) = \|x^* - x\|_A = \left((x^* - x)^T A (x^* - x)\right)^{1/2} = (A(x^* - x), x^* - x)^{1/2},$$

Proof. For $\tilde{x}$ to be the minimizer of $E(\tilde{x})$, it is necessary and sufficient that $x^* - \tilde{x}$ be A-orthogonal to all vectors in $\mathcal{K}$, i.e., iff

$$(x^* - \tilde{x})^T A v = 0, \quad \forall v \in \mathcal{K},$$

or equivalently

$$(b - A\tilde{x})^T v = 0, \quad \forall v \in \mathcal{K},$$

which is the Galerkin condition.

Orthogonal Projection Methods

Given the Krylov subspace $\mathcal{K}_m(A, r_0)$, we seek an approximate solution x_m from $x_0 + \mathcal{K}_m(A, r_0)$ such that

$$b - Ax_m \perp \mathcal{K}_m(A, r_0) \tag{1}$$

By Arnoldi process and given $q_1 = r_0/\beta, \beta = \|r_0\|$, we obtain:

$$\begin{aligned} Q_m^T A Q_m &= H_m &(2) \\ Q_m^T r_0 &= Q_m^T(\beta q_1) = \beta e_1 &(3) \end{aligned}$$

The approximate solution is obtained as

$$\begin{aligned} x_m = x_0 + Q_m y_m, \text{ where} &\quad (4) \\ y_m = H_m^{-1}(\beta e_1) &\quad (5) \end{aligned}$$

since $0 = Q_m^T r_m = Q_m^T(b - Ax_m) = Q_m^T r_0 - Q_m^T A Q_m y_m$

Symmetric matrices and Lanczos algorithm

Given the Arnoldi process, $A \in \mathbb{R}^{n \times n}$, $\bar{H}_m$ and $H_m \in \mathbb{R}^{m \times m}$ formed from $\bar{H}_m$ by deleting its last row, we have:

$$Q_m^T A Q_m = H_m \tag{6}$$

- If A is symmetric, then H_m is tridiagonal, this leads to:
  - Lanczos process: mainly used for eigenvalue computation
  - Conjugate Gradient: most used algorithm for solving linear systems when A is symmetric positive definite

Symmetric Lanczos algorithm

Let $\alpha_j \equiv h_{jj} \quad \beta_j \equiv h_{j-1,j}$,

$$H_m = \begin{pmatrix} \alpha_1 & \beta_2 & & & \\ \beta_2 & \alpha_2 & \beta_3 & & \\ & . & . & . & \\ & & \beta_{m-1} & \alpha_{m-1} & \beta_m \\ & & & \beta_m & \alpha_m \end{pmatrix} \tag{7}$$

Symmetric Lanczos uses a three-term recurrence:

$$\beta_{j+1} q_{j+1} = A q_j - \alpha_j q_j - \beta_j q_{j-1},$$

Algorithm 1 Symmetric Lanczos for linear systems

$r_0 = b - Ax_0, \beta = \|r_0\|_2,\ q_1 = r_0/\beta$
for $j = 1 : m$ **do**
  $w_{j+1} = Aq_j - \beta_j q_{j-1}$ (If $j = 1$ set $\beta_1 q_0 \equiv 0$)
  $\alpha_j = \langle w_{j+1}, q_j \rangle$
  $w_{j+1} = w_{j+1} - \alpha_j q_j$
  $\beta_{j+1} = \|w_{j+1}\|_2$. If $\beta_{j+1} = 0$ then Stop
  $q_{j+1} = w_{j+1}/\beta_{j+1}$
end for
Set $H_m = tridiag(\beta_j, \alpha_j, \beta_{j+1})$ and $Q_m = [q_1, \ldots, q_m]$
Compute $y_m = H_m^{-1}(\beta e_1)$ and $x_m = x_0 + Q_m y_m$

Conjugate gradient (Hestenes, Stieffel, 52)

- A Krylov projection method for SPD matrices where $\mathscr{L}_k = \mathcal{K}_k(A, r_0)$.
- Finds $x^* = A^{-1}b$ by minimizing the quadratic function

$$\begin{aligned} \phi(x) &= \frac{1}{2}(x)^t Ax - b^t x \\ \triangledown\phi(x) &= Ax - b = 0 \end{aligned}$$

- After j iterations of CG,

$$||x^* - x_j||_A \leq 2||x^* - x_0||_A \left(\frac{\sqrt{\kappa(A)} - 1}{\sqrt{\kappa(A)} + 1} \right)^j, \tag{8}$$

where x_0 is starting vector, $||x||_A = \sqrt{x^T Ax}$ and $\kappa(A) = |\lambda_{max}(A)|/|\lambda_{min}(A)|$.

Conjugate gradient

- Computes A-orthogonal search directions by conjugation of the residuals

$$\begin{cases} p_1 &= r_0 = -\bigtriangledown \phi(x_0) \\ p_k &= r_{k-1} + \beta_k p_{k-1} \end{cases} \tag{9}$$

- At k-th iteration,

$$p_k = r_{k-1} + \beta_k p_{k-1} \tag{10}$$

$$x_k = x_{k-1} + \alpha_k p_k = argmin_{x \in x_0 + \mathcal{K}_k(A, r_0)} \phi(x) \tag{11}$$

$$r_k = r_{k-1} - \alpha_k A p_k \tag{12}$$

where α_k is the step along p_k.

- CG algorithm obtained by imposing the orthogonality and the conjugacy conditions

$$\begin{aligned} r_k^T r_i &= 0, \text{for all } i \neq k, \\ p_k^T A p_i &= 0, \text{for all } i \neq k. \end{aligned}$$

CG derivation

Since we have $x_k = x_{k-1} + \alpha_k p_k$ we obtain

$$r_k = r_{k-1} - \alpha_k A p_k \text{ and } (r_k, r_{k-1}) = 0 \text{ hence}$$
$$r_{k-1}^T r_{k-1} - \alpha_k r_{k-1}^T A p_k = 0 \implies \alpha_k = \frac{(r_{k-1}, r_{k-1})}{(Ap_k, r_{k-1})}$$

Since we have $p_k = r_{k-1} + \beta_k p_{k-1}$,

$$(Ap_k, r_{k-1}) = (Ap_k, p_k - \beta_k p_{k-1}) = (Ap_k, p_k) \implies \alpha_k = \frac{(r_{k-1}, r_{k-1})}{(Ap_k, p_k)}$$

Since $p_k = r_{k-1} + \beta_k p_{k-1}$ is A-orthogonal to p_{k-1} we obtain

$$\beta_k = -\frac{(r_{k-1}, Ap_{k-1})}{(p_{k-1}, Ap_{k-1})} \text{ and } Ap_{k-1} = \frac{1}{\alpha_{k-1}}(r_{k-2} - r_{k-1}) \implies \beta_k = \frac{(r_{k-1}, r_{k-1})}{(r_{k-2}, r_{k-2})}$$

CG algorithm

Algorithm 2 The CG Algorithm

$r_0 = b - Ax_0,\ \rho_0 = ||r_0||_2^2,\ p_1 = r_0,\ k = 1$
while ($\sqrt{\rho_k} > \epsilon ||b||_2$ and $k < k_{max}$) **do**
  if $(k \neq 1)$ **then**
    $\beta_k = (r_{k-1}, r_{k-1})/(r_{k-2}, r_{k-2})$
    $p_k = r_{k-1} + \beta_k p_{k-1}$
  end if
  $\alpha_k = (r_{k-1}, r_{k-1})/(Ap_k, p_k)$
  $x_k = x_{k-1} + \alpha_k p_k$
  $r_k = r_{k-1} - \alpha_k A p_k$
  $\rho_k = ||r_k||_2^2$
  $k = k + 1$
end while

Properties of CG

- The directions $p_1, \dots p_k$ are A-conjugate, the following properties are satisfied:

$$\begin{aligned}(Ap_k, p_j) &= 0, \text{ for all } k, j, k \neq j \\ (r_k, r_j) &= 0, \text{ for all } k, j, k \neq j \\ (p_k, r_j) &= 0, \text{ for all } k, j, k < j\end{aligned}$$

- The Krylov subspace is spanned by the residuals and the search directions:

$$\mathcal{K}_k(A, r_0) = \text{span}\{r_0, r_1, ..., r_{k-1}\} = \text{span}\{p_0, p_1, ..., p_{k-1}\}$$

Adviced exercice: prove the above relations, e.g. by using recurrence on equations (10), (11), (12).
We do not prove (11) and (8), the proofs can be found in [Saad, 2003]

Plan

Krylov subspace methods

Enlarged Krylov methods

Definition and properties

Numerical and parallel performance results

Enlarged Krylov methods [LG, Moufawad, Nataf, 14]

- Partition the matrix into N domains
- Split the residual r_0 into t vectors corresponding to the N domains, obtain R_0^e,

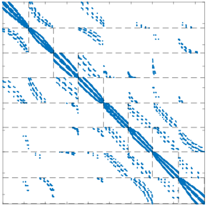

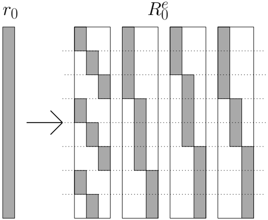

- Generate t new basis vectors, obtain an **enlarged** Krylov subspace

$$\mathcal{K}_{t,k}(A, r_0) = \text{span}\{R_0^e, AR_0^e, A^2 R_0^e, ..., A^{k-1} R_0^e\}$$

- Search for the solution of the system $Ax = b$ in $\mathcal{K}_{t,k}(A, r_0)$

Properties of enlarged Krylov subspaces

- The Krylov subspace $\mathcal{K}_k(A, r_0)$ is a subset of the enlarged one

$$\mathcal{K}_k(A, r_0) \subset \mathcal{K}_{t,k}(A, r_0)$$

- For all $k < k_{max}$ the dimensions of $\mathcal{K}_{t,k}$ and $\mathcal{K}_{t,k+1}$ are stricltly increasing by some number i_k and i_{k+1} respectively, where

$$t \geq i_k \geq i_{k+1} \geq 1.$$

- The enlarged subspaces are increasing subspaces, yet bounded.

$$\mathcal{K}_{t,1}(A, r_0) \subsetneq ... \subsetneq \mathcal{K}_{t,k_{max}-1}(A, r_0) \subsetneq \mathcal{K}_{t,k_{max}}(A, r_0) = \mathcal{K}_{t,k_{max}+q}(A, r_0), \forall q > 0$$

- The solution of the system $Ax = b$ belongs to the subspace $x_0 + \mathcal{K}_{t,k_{max}}$.

Enlarged Krylov subspace methods based on CG

Defined by the subspace $\mathcal{K}_{t,k}$ and the following two conditions:

1. Subspace condition: $x_k \in x_0 + \mathcal{K}_{t,k}$
2. Orthogonality condition: $r_k \perp \mathcal{K}_{t,k}$

- At each iteration, the new approximate solution x_k is found by minimizing $\phi(x) = \frac{1}{2}(x)^t Ax - b^t x$ over $x_0 + \mathcal{K}_{t,k}$:

$$\phi(x_k) = min\{\phi(x), \forall x \in x_0 + \mathcal{K}_{t,k}(A, r_0)\}$$

- Can be seen as a particular case of a block Krylov method
  - $AX = S(b)$, such that $S(b)ones(t,1) = b$; $R_0^e = AX_0 - S(b)$
  - Orthogonality condition involves the block residual $R_k \perp \mathcal{K}_{t,k}$

Enlarged Krylov subspace methods based on CG

Defined by the subspace $\mathcal{K}_{t,k}$ and the following two conditions:

1. Subspace condition: $x_k \in x_0 + \mathcal{K}_{t,k}$
2. Orthogonality condition: $r_k \perp \mathcal{K}_{t,k}$

- At each iteration, the new approximate solution x_k is found by minimizing $\phi(x) = \frac{1}{2}(x)^t Ax - b^t x$ over $x_0 + \mathcal{K}_{t,k}$:

$$\phi(x_k) = min\{\phi(x), \forall x \in x_0 + \mathcal{K}_{t,k}(A, r_0)\}$$

- Can be seen as a particular case of a block Krylov method
  - $AX = S(b)$, such that $S(b)ones(t, 1) = b$; $R_0^e = AX_0 - S(b)$
  - Orthogonality condition involves the block residual $R_k \perp \mathcal{K}_{t,k}$

Convergence analysis

Given

- A is an SPD matrix, x^* is the solution of $Ax = b$
- $||x^* - \overline{x}_k||_A$ is the k^{th} error of CG
- $||x^* - x_k||_A$ is the k^{th} error of ECG

Result

CG	**ECG**
$\|\|x^* - \overline{x}_k\|\|_A \leq 2\|\|x^* - x_0\|\|_A \left(\frac{\sqrt{\kappa}-1}{\sqrt{\kappa}+1}\right)^k$	$\|\|x^* - x_k\|\|_A \leq C\|\|x^* - x_0\|\|_A \left(\frac{\sqrt{\kappa_t}-1}{\sqrt{\kappa_t}+1}\right)^k$
where $\kappa = \frac{\lambda_{max}(A)}{\lambda_{min}(A)}$	where $\kappa_t = \frac{\lambda_{max}(A)}{\lambda_t(A)}$ C is a const indpdt. of k, dpdt. of t

Proof of convergence of ECG can be found in [Grigori and Tissot, 2019].

Classical CG *vs.* Enlarged CG derived from Block CG

Algorithm 3 Classical CG

1: $p_1 = r_0(r_0^\top A r_0)^{-1/2}$
2: **while** $||r_{k-1}||_2 > \varepsilon||b||_2$ **do**
3: $\quad \alpha_k = p_k^\top r_{k-1}$
4: $\quad x_k = x_{k-1} + p_k \alpha_k$
5: $\quad r_k = r_{k-1} - A p_k \alpha_k$
6: $\quad z_{k+1} = r_k - p_k(p_k^\top A r_k)$
7: $\quad p_{k+1} = z_{k+1}(z_{k+1}^\top A z_{k+1})^{-1/2}$
8: **end while**

Cost per iteration

$\#$ flops $= O(\frac{n}{P}) \leftarrow$ BLAS 1 & 2
$\#$ words $= O(1)$
$\#$ messages $= O(1)$ from SpMV + $O(logP)$ from dot prod + norm

Algorithm 4 ECG

1: $P_1 = R_0^e(R_0^{e\top} A R_0^e)^{-1/2}$
2: **while** $|| \sum_{i=1}^{\top} R_k^{(i)} ||_2 < \varepsilon||b||_2$ **do**
3: $\quad \alpha_k = P_k^\top R_{k-1}$ $\qquad \triangleright$ $t \times t$ matrix
4: $\quad X_k = X_{k-1} + P_k \alpha_k$
5: $\quad R_k = R_{k-1} - A P_k \alpha_k$
6: $\quad$ Construct Z_{k+1} s.t. $Z_{k+1}^\top A P_i = 0, \forall i \le k$
7: $\quad P_{k+1} = Z_{k+1}(Z_{k+1}^\top A Z_{k+1})^{-1/2}$
8: **end while**
9: $x = \sum_{i=1}^{\top} X_k^{(i)}$

Cost per iteration

$\#$ flops $= O(\frac{nt^2}{P}) \leftarrow$ BLAS 3
$\#$ words $= O(t^2) \leftarrow$ Fit in the buffer
$\#$ messages $= O(1)$ from SpMV + $O(logP)$ from A-ortho

Test cases: boundary value problem

2D and 3D Skyscraper Problem - SKY2D,3D

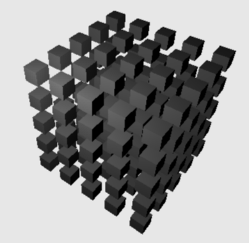

$$
\begin{aligned}
-div(\kappa(x)\nabla u) &= f \ in \ \Omega \\
u &= 0 \ on \ \partial\Omega_D \\
\frac{\partial u}{\partial n} &= 0 \ on \ \partial\Omega_N
\end{aligned}
$$

discretized on a 3D grid , where

$$
\kappa(x) = \begin{cases} 10^3 * ([10 * x_2] + 1), if \ [10 * x_i] = 0 mod(2), \ i = 1, 2, 3, \\ 1, \qquad\qquad\qquad\qquad\qquad\qquad otherwise. \end{cases}
$$

Test cases (contd)

Linear elasticity 3D problem

$$\begin{aligned}
\operatorname{div}(\sigma(u)) + f &= 0 && \text{on } \Omega,\\
u &= u_D && \text{on } \partial\Omega_D,\\
\sigma(u)\cdot n &= g && \text{on } \partial\Omega_N,
\end{aligned}$$

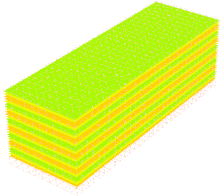

Figure: The distribution of Young's modulus

- $u \in \mathbb{R}^d$ is the unknown displacement field, f is some body force.
- Young's modulus E and Poisson's ratio ν take two values, $(E_1, \nu_1) = (2 \cdot 10^{11}, 0.25)$, and $(E_2, \nu_2) = (10^7, 0.45)$.
- Cauchy stress tensor $\sigma(u)$ is given by Hooke's law, defined by E and ν.

Matrices Generated with FreeFem++ (F. Hecht, Sorbonne Université)
Linear Elasticity discretized using $\mathbb{P}_1$ FE, $1600 \times Y \times Y$ grid

Enlarged CG: numerical results

- Block Jacobi preconditioner (1024 blocks)
- Stopping criterion 10^{-6}, initial block size 32
- BRRHS-CG: block method with $t-1$ random rhs

matrix	$n(A)$	$nnz(A)$
SKY2D	10,000	49,600
Ela3D100	36,663	1,231,497
Ela2D200	80,802	964,800

		PCG	BRRHS-CG		ECG	
	red. size	iter	iter	$\dim(\mathcal{K}_{32,k})$	iter	$\dim(\mathcal{K}_{32,k})$
SKY2D	$\times$	655	61	1952	57	1824
	$\checkmark$	655	61	1739	59	1546
Ela3D100	$\times$	955	102	3264	109	3488
	$\checkmark$	955	102	3093	116	2384
Ela2D200	$\times$	4551	255	8160	253	8096
	$\checkmark$	4551	258	7331	266	6553

Enlarged CG: parallel performance

- Stopping criterion 10^{-5}, blocks Jacobi $= \#$MPI
- Performance study on:
  – Kebnekaise (Suede), Intel Xeon (Broadwell), 28 MPI process/node
  – Cori NERSC, Intel KNL, 68 cores each

	ECG($\mathcal{K}_{24}$)		CG	
# MPI	# iter	res	# iter	res
252	513	1.3E-4	13,626	1.3E-4
504	531	1.9E-4	15,819	1.9E-4
1,008	606	2.6E-4	17,023	2.7E-4
2,016	696	2.6E-4	19,047	2.7E-4

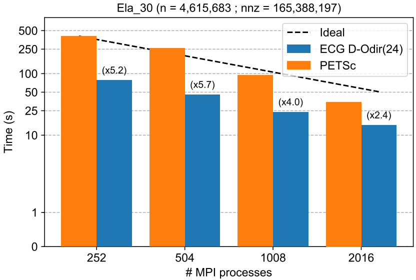

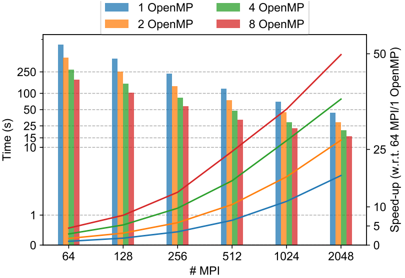

References (1)

Ashby, S. F., Manteuffel, T. A., and Saylor, P. E. (1990).
A taxonomy for conjugate gradient methods.
SIAM Journal on Numerical Analysis, 27(6):1542–1568.

Chan, T. F. and Mathew, T. P. (1994).
Domain decomposition algorithms.
Acta Numerica, 3:61–143.

Dolean, V., Jolivet, P., and Nataf, F. (2015).
An introduction to domain decomposition methods.
Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA.
Algorithms, theory, and parallel implementation.

Grigori, L., Moufawad, S., and Nataf, F. (2016).
Enlarged Krylov Subspace Conjugate Gradient Methods for Reducing Communicatio n.
SIAM Journal on Scientific Computing, 37(2):744–773.
Also as INRIA TR 8266.

Grigori, L. and Tissot, O. (2019).
Scalable linear solvers based on enlarged krylov subspaces with dynamic reduct ion of search directions.
SIAM Journal on Scientific Computing, 34(1):206–239.

Hestenes, M. R. and Stiefel., E. (1952).
Methods of conjugate gradients for solving linear systems.
Journal of research of the National Bureau of Standards., 49:409–436.

O'Leary., D. P. (1980).
The block conjugate gradient algorithm and related methods.
Linear Algebra and Its Applications, 29:293–322.

References (2)

Saad, Y. (2003).
Iterative Methods for Sparse Linear Systems.
Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2nd edition.

Tang, J. M., Nabben, R., Vuik, C., and Erlangga, Y. A. (2009).
Comparison of two-level preconditioners derived from deflation, domain decomposition and multigrid methods.
J. Sci. Comput., 39:340–370.