





Parallel Performance and Numerical Stability of Communication Hiding Pipelined Krylov Subspace Methods

10th International Workshop on Parallel Matrix Algorithms and Applications (PMAA'18) ${\rm ETH}\ {\rm Z\ddot{u}rich},\ {\rm June}\ 27\text{-}29,\ 2018$

University of Antwerp* [Belgium], LBNL[‡] [USA], INRIA Bordeaux[†] [France]

S. Cools*, J. Cornelis*, W. Vanroose*, T. Mijieux*, P. Ghysels[‡], E. F. Yetkin[†], E. Agullo[†], L. Giraud[†]

* E-mail: siegfried.cools@uantwerp.be

Universiteit Antwerpen



Motivation Exascale systems projection

	Today's Systems	Predicted Exascale Systems*	Factor Improvement	
System Peak	10 ¹⁶ flops/s	10^{18} flops/s	100	
Node Memory Bandwidth	10^2 GB/s	10^3 GB/s	10	
Interconnect Bandwidth	$10^1 \mathrm{GB/s}$	10 ² GB/s	10	
Memory Latency	10^{-7} s	$5\cdot 10^{-8}$ s	2	
Interconnect Latency	$10^{-6} s$	$5\cdot 10^{-7}$ s	2	

*Sources: from P. Beckman (ANL), J. Shalf (LBL), and D. Unat (LBL)

- Data movement (communication) is much more expensive than flops (computation) in terms of both time and energy
- Reducing time spent moving data/waiting for data will be essential for exascale applications

 \Rightarrow Communication avoiding / Communication hiding

Motivation Communication hiding vs. communication avoiding

Communication cost has motivated several approaches to reducing global synchronization cost in Krylov subspace methods:

Avoiding communication: *s*-step Krylov subspace methods * [A. Chronopoulous, J. Demmel, M. Hoemmen, E. Carson, L. Grigori, J. Erhel, ...]

- Compute iterations in blocks of s (change of Krylov subspace basis)
- Reduces number of synchronizations per iteration by a factor of $\mathcal{O}(s)$

Hiding communication: Pipelined Krylov subspace methods * [P. Ghysels, W. Vanroose, S. C., P. Sanan, B. Gropp, I. Yamazaki, ...]

- Introduce auxiliary vectors to decouple SpMV and inner products
- Enables overlapping of communication and computation

* Both equivalent to corresponding Krylov subspace methods in exact arithmetic

Motivation



Communication hiding vs. communication avoiding

June 27, 2018 Communication co g global 10:15 AM - 12:15 PM synchronization co 1.2.B: Scalable communication-reducing Krylov subspace methods Avoiding comm hods * Organizers: S. Cools [A. Chronopoulous, J. Location: CAB G61 10:15 AM Jeffrey Cornelis: Hiding global communication in the e basis) Compute iter Conjugate Gradient method using deep pipelines • 10:45 AM Hussam Al Daas: Recycling Krylov method for the solution of Reduces num of $\mathcal{O}(s)$ sequence of linear systems 11:15 AM Sicong Zhuang: Iteration-Fusing Conjugate Gradient 11:45 AM Emmanuel Agullo: Partial convergence in block Krylov solvers thods * Hiding commur [P. Ghysels, W. Vanrod Introduce aux lucts Enables overl Both equivalent to d Thumbs up to all MS speakers!



Krylov subspace methods General concepts

Iteratively improve an approximate solution of linear system Ax = b,

$$x_i \in x_0 + \mathcal{K}_i(A, r_0) = x_0 + \operatorname{span}\{r_0, Ar_0, A^2r_0, \dots, A^{i-1}r_0\}$$



Krylov subspace methods General concepts

Iteratively improve an approximate solution of linear system Ax = b,

 $x_i \in x_0 + \mathcal{K}_i(A, r_0) = x_0 + \operatorname{span}\{r_0, Ar_0, A^2r_0, \dots, A^{i-1}r_0\}$

- minimize certain error measure over Krylov subspace K_i(A, r₀)
- Krylov subspace methods:
 Conjugate Gradients (CG),
 Lanczos, GMRES, MinRES,
 BiCG, CGS, BiCGStab, CGLS, ...
- Preconditioners:

AMG & GMG, Domain Decomposition Methods, FETI, BDDC, Incomplete factorization, Physics based preconditioners, ...



Krylov subspace methods General concepts

Iteratively improve an approximate solution of linear system Ax = b,

 $x_i \in x_0 + \mathcal{K}_i(A, r_0) = x_0 + \operatorname{span}\{r_0, Ar_0, A^2r_0, \dots, A^{i-1}r_0\}$

- ► minimize certain error measure over Krylov subspace K_i(A, r₀)
- Krylov subspace methods:
 Conjugate Gradients (CG),
 Lanczos, GMRES, MinRES,
 BiCG, CGS, BiCGStab, CGLS, ...
- Preconditioners:

AMG & GMG, Domain Decomposition Methods, FETI, BDDC, Incomplete factorization, Physics based preconditioners, ...

- usually in combination with sparse linear algebra/stencil application
- three algorithmic building blocks:
 - i. dot-product
 - $\mathcal{O}(N)$ flops
 - global synchronization (MPI_Allreduce)
 - ii. SpMV
 - $\mathcal{O}(nnz)$ flops
 - neighbor communication only
 - iii. axpy
 - $\mathcal{O}(N)$ flops
 - no communication



Krylov subspace methods Classical CG

Algorithm CG 1: procedure $CG(A, b, x_0)$ 2: $r_0 := b - Ax_0; p_0 = r_0$ 3: for i = 0, ... do 4: $s_i := Ap_i$ 5: $\alpha_i := (r_i, r_i) / (s_i, p_i)$ 6: $x_{i+1} := x_i + \alpha_i p_i$ 7: $r_{i+1} := r_i - \alpha_i s_i$ 8: $\beta_{i+1} := (r_{i+1}, r_{i+1}) / (r_i, r_i)$ 9: $p_{i+1} := r_{i+1} := r_{i+1} p_{i+1}$

10: end for 11: end procedure

Hestenes & Stiefel (1952)

i. dot-products

- ▶ 2 global reductions: latency dominated
- scales as log₂(#partitions)

ii. SpMV

- computationally expensive
- good scaling (minor comm.)

iii. axpy's

- vector operations (recurrence relations)
- perfect scaling (no comm.)



dot-pr

SpMV

axpy



Krylov subspace methods Global reduction latency



T. Hoeffler, T. Schneider and A. Lumsdaine, SC10, 2010



Algorithm Pipelined CG	
1: procedure PIPE-CG (A, b, x_0)	
2: $r_0 := b - Ax_0; w_0 := Ar_0$	
3: for $i = 0,$ do	
4: $\gamma_i := (r_i, r_i)$	
5: $\delta := (w_i, r_i)$	
$6: - q_i := Aw_i$	
7: if $i > 0$ then	
8: $\beta_i := \gamma_i / \gamma_{i-1}; \alpha_i := (\delta / \gamma_i - \beta_i / \epsilon)$	$(x_{i-1})^{-1}$
9: else	
10: $\beta_i := 0; \alpha := \gamma_i / \delta$	dot-pr
11: end if	Crob(1)/
12: $z_i := q_i + \beta_i z_{i-1}$	Shinin
13: $s_i := w_i + \beta_i s_{i-1}$	ахру
14: $p_i := r_i + \beta_i p_{i-1}$	
15: $x_{i+1} := x_i + \alpha_i p_i$	
16: $r_{i+1} := r_i - \alpha_i s_i$	
17: $w_{i+1} := w_i - \alpha_i z_i$	
18: end for	
19: end procedure	

Pipelined CG = re-engineered version of CG for improved parallel performance

Ghysels & Vanroose (2014)



Algorithm Pipelined CG	
1: procedure PIPE-CG (A, b, x_0)	
2: $r_0 := b - Ax_0; w_0 := Ar_0$	
3: for $i = 0,$ do	
4: $\gamma_i := (r_i, r_i)$	
5: $\delta := (w_i, r_i)$	
$6: - q_i := Aw_i$	
7: if $i > 0$ then	
8: $\beta_i := \gamma_i / \gamma_{i-1}; \ \alpha_i := (\delta / \gamma_i - \beta_i)$	$\beta_i/\alpha_{i-1})^{-1}$
9: else	
10: $\beta_i := 0; \alpha := \gamma_i / \delta$	dot-pr
11: end if	CroMAV
12: $z_i := q_i + \beta_i z_{i-1}$	Shinin
13: $s_i := w_i + \beta_i s_{i-1}$	ахру
14: $p_i := r_i + \beta_i p_{i-1}$	
15: $x_{i+1} := x_i + \alpha_i p_i$	
16: $r_{i+1} := r_i - \alpha_i s_i$	
17: $w_{i+1} := w_i - \alpha_i z_i$	
18: end for	
19: end procedure	

Pipelined CG = re-engineered version of CG for improved parallel performance

► Introduce new auxiliary variables ...
s_i = Ap_i, w_i = Ar_i, z_i = As_i = A²p_i

Ghysels & Vanroose (2014)



Algorithm Pipelined CG	
1: procedure PIPE-CG (A, b, x_0)	
2: $r_0 := b - Ax_0; w_0 := Ar_0$	
3: for $i = 0,$ do	
4: $\gamma_i := (r_i, r_i)$	
5: $\delta := (w_i, r_i)$	
$6: - q_i := Aw_i$	
7: if $i > 0$ then	
8: $\beta_i := \gamma_i / \gamma_{i-1}; \ \alpha_i := (\delta / \gamma_i - \beta_i)$	$\beta_i/\alpha_{i-1})^{-1}$
9: else	
10: $\beta_i := 0; \alpha := \gamma_i / \delta$	dot-pr
11: end if	CroM/V
12: $z_i := q_i + \beta_i z_{i-1}$	Shinin
13: $s_i := w_i + \beta_i s_{i-1}$	ахру
14: $p_i := r_i + \beta_i p_{i-1}$	
15: $x_{i+1} := x_i + \alpha_i p_i$	
16: $r_{i+1} := r_i - \alpha_i s_i$	
17: $w_{i+1} := w_i - \alpha_i z_i$	
18: end for	
19: end procedure	

Ghysels & Vanroose (2014)

Pipelined CG = re-engineered version of CG for improved parallel performance

Introduce new auxiliary variables ...
 s_i = Ap_i, w_i = Ar_i, z_i = As_i = A²p_i

that are computed recursively, e.g.:

 $p_i = r_i + \beta_i p_{i-1}$ $\Rightarrow s_i = A p_i = w_i + \beta_i s_{i-1}$ with $w_i := A r_i$



Algorithm Pipelined CG 1: procedure PIPE-CG(A, b, x_0) $r_0 := b - Ax_0; w_0 := Ar_0$ 2: for $i = 0, \ldots$ do 3: 4: $\gamma_i := (r_i, r_i)$ 5: $\delta := (w_i, r_i)$ $q_i := Aw_i$ 6: -7: if i > 0 then $\beta_i := \gamma_i / \gamma_{i-1}; \alpha_i := (\delta / \gamma_i - \beta_i / \alpha_{i-1})^{-1}$ 8: ٩. else 10. $\beta_i := 0; \alpha := \gamma_i / \delta$ dot-pr end if 11. **SpMV** 12: _____ $z_i := q_i + \beta_i z_{i-1}$ 13: $s_i := w_i + \beta_i s_{i-1}$ 14: $p_i := r_i + \beta_i p_{i-1}$ axpy 15: _____ $x_{i+1} := x_i + \alpha_i p_i$ 16: $r_{i+1} := r_i - \alpha_i s_i$ 17: $w_{i+1} := w_i - \alpha_i z_i$ end for 18: 19: end procedure

Ghysels & Vanroose (2014)

- i. Communication avoiding: dot-products grouped in one global reduction phase per iteration
- ii. Communication hiding: overlap global synchronization with SpMV (+ Prec) computation
- iii. No free lunch: Additional recurrence relations, i.e. axpy's, for auxiliary var's s_i = Ap_i, w_i = Ar_i, z_i = As_i





Krylov subspace methods Generalization: deep pipelines

Classic KSM:



Pipelined KSM:



Deep pipelined KSM:





Classic KSM:



Pipelined KSM:



Deep pipelined KSM:



Krylov subspace methods Generalization: deep pipelines

Consider the Arnoldi relation

 $AV_i = V_{i+1}H_{i+1,i}$

with V_i the Krylov subspace basis and $H_{i+1,i}$ upper Hessenberg. Introduce the auxiliary vectors $Z_{i+1} = [z_0, z_1, \dots, z_i]$ as

$$z_j := \begin{cases} v_0 & j = 0, \\ P_j(A)v_0 & 0 < j \le I, \\ P_l(A)v_{j-l} & j > l, \end{cases}$$

with polynomials $P_l(t)$ of fixed order l, where l is the pipeline length

$$P_l(t) := \prod_{j=0}^{l-1} (t - \sigma_j).$$

Krylov subspace methods Pipelined KSM with deep pipeline

Writing the Arnoldi relation $AV_i = V_{i+1}H_{i+1,i}$ for basis vector v_{j-1} and applying $P_i(A)$ to both sides, we obtain recurrence relations for z_j :

$$v_{j-l} = \frac{Av_{j-l-1} - \sum_{k=0}^{j-l-1} h_{k,j-l-1}v_k}{h_{j-l,j-l-1}} \quad \Rightarrow \quad \mathbf{z}_j = \frac{A\mathbf{z}_{j-1} - \sum_{k=0}^{j-l-1} h_{k,j-l-1}\mathbf{z}_{k+l}}{h_{j-l,j-l-1}}.$$

This implies an Arnoldi-like relation for the auxiliary variables

$$AZ_i = Z_{i+1}B_{i+1,i},$$

with $B_{i+1,i}$ the shifted Hessenberg matrix.

<u>Basis transformation</u>: Z_i and V_i both span the *i*-th Krylov subspace, thus there an upper triangular transformation matrix G_i such that

$$Z_i = V_i G_i$$
.

▶ Recursive calculation of entries of G_i ($\forall j = 0, 1, ..., k \leq i$):

$$g_{j,k} = (z_k, v_j) = \left(z_k, rac{z_j - \sum_{m=0}^{j-1} g_{m,j} v_m}{g_{j,j}}
ight) = rac{(z_k, z_j) - \sum_{m=0}^{j-1} g_{m,j} g_{m,k}}{g_{j,j}}$$

▶ Recursive calculation of the Hessenberg matrix $H_{i+1,i}$:

$$\begin{aligned} H_{k+1,k} &= V_{k+1}^{T} A V_{k} \\ &= V_{k+1}^{T} A Z_{k} G_{k}^{-1} \quad \text{using} \quad Z_{k} = V_{k} G_{k} \\ &= V_{k+1}^{T} Z_{k+1} B_{k+1,k} G_{k}^{-1} \quad \text{using} \quad A Z_{k} = Z_{k+1} B_{k+1,k} \\ &= G_{k+1} B_{k+1,k} G_{k}^{-1} \quad \text{using} \quad G_{k+1} = V_{k+1}^{T} Z_{k+1} \\ &= \begin{pmatrix} H_{k,k-1} & (G_{k} b_{:,k} + g_{:,k+1} b_{k+1,k} - H_{k,k-1} g_{:,k}) g_{k,k}^{-1} \\ 0 & g_{k+1,k+1} b_{k+1,k} g_{k,k}^{-1} \end{pmatrix} \\ &\text{using} \quad H_{k,k-1} G_{k-1} = G_{k} B_{k,k-1} \end{aligned}$$

► Combining all relations leads to the deep pipelined p(*I*)-GMRES algorithm.
■ Ghysels & Vanroose (2013)



 $\label{eq:Krylov subspace methods} \ensuremath{\mathsf{Pipelined CG}} \ensuremath{\mathsf{with deep pipeline}} \ensuremath{\mathsf{Pipeline}}$

In case of an SPD matrix A the following features simplify the algorithm:

- ► *H*_{*i*+1,*i*} is tridiagonal
- G_i has a (2l + 1)-nonzero diagonal band structure.



The solution x_i can be constructed using a recursively computed search direction p_i, instead of using the entire Krylov basis V_i:

$$p_j = (v_j - \delta_{j-1}p_{j-1})/\eta_j, \qquad x_j = x_{j-1} + \xi_{j-1}p_{j-1}.$$

Liesen & Strakos, Krylov Subspace Methods (2012)



In case of an SPD matrix A the following features simplify the algorithm:

- ► *H*_{*i*+1,*i*} is tridiagonal
- G_i has a (2l + 1)-nonzero diagonal band structure.



The solution x_i can be constructed using a recursively computed search direction p_i, instead of using the entire Krylov basis V_i:

$$p_j = (v_j - \delta_{j-1}p_{j-1})/\eta_j, \qquad x_j = x_{j-1} + \xi_{j-1}p_{j-1}.$$

Liesen & Strakos, Krylov Subspace Methods (2012)



Algorithm Deep pipelined Conjugate Gradients $(p(\ell)$ -CG)

1: $r_0 := b - Ax_0;$ 2: $v_0 := r_0 / \|r_0\|_2$; 3: $z_0 := v_0; \ g_{0,0} := 1;$ 4: for i = 0, ..., m + l do $z_{i+1} := \begin{cases} (A - \sigma_i I) z_i, & i < l \\ A z_i, & i > l \end{cases}$ 5 6. a := i - l7. if a > 0 then 8: $g_{j,a+1} := (g_{j,a+1} - \sum_{k=a+1-2l}^{j-1} g_{k,j}g_{k,a+1})/g_{j,j}; \qquad j = a - l + 2, \dots, a$ $g_{a+1,a+1} := \sqrt{g_{a+1,a+1} - \sum_{k=a+1-2l}^{a} g_{k,a+1}^2};$ 9: # Check for breakdown and restart if required 10: 11: if a < l then $\gamma_a := (q_{a,a+1} + \sigma_a q_{a,a} - \delta_{a-1} q_{a-1,a})/q_{a,a};$ 12:13: $\delta_a := q_{a+1,a+1}/q_{a,a};$ 14: else $\gamma_a := (g_{a,a}\gamma_{a-l} + g_{a,a+1}\delta_{a-l} - \delta_{a-1}g_{a-1,a})/g_{a,a};$ 15: $\delta_a := (q_{a+1,a+1}\delta_{a-1})/q_{a,a};$ 16: end if 17: $v_{a+1} := (z_{a+1} - \sum_{j=a-2l+1}^{a} g_{j,a+1}v_j)/g_{a+1,a+1};$ 18: 19: $z_{i+1} := (z_{i+1} - \gamma_a z_i - \delta_{a-1} z_{i-1})/\delta_a;$ end if 20:21:if a < 0 then $g_{j,i+1} := (z_{i+1}, z_j); \qquad j = 0, \dots, i+1$ 22: $23 \cdot$ else $g_{j,i+1} := \begin{cases} (z_{i+1}, v_j); & j = \max(0, i-2l+1), \dots, a+1\\ (z_{i+1}, z_j); & j = a+2, \dots, i+1 \end{cases}$ 24:25:end if

Cornelis et al. (2018)





Cornelis et al. (2018)





Cornelis et al. (2018)

Krylov subspace methods Parallel performance of pipelined CG

Strong scaling experiments - PETSc 3.6.3/3.7.6 library - MPICH 3.1.3/3.3a2



Krylov subspace methods Overview of pipelined Krylov methods

Pipelined GMRES

4

- depth-one pipeline: p-GMRES
- deep pipelines: $p(\ell)$ -GMRES: compute ℓ new Krylov subspace basis vectors (SpMV's) during global communication and orthogonalize after ℓ iterations.

$$Z_{i+1} = [z_0, z_1, \dots, z_{i-\ell}, \underbrace{z_{i-\ell+1}, \dots, z_i}_{\ell}]$$

 V_{i} and - [16 16 V_{i}]

- Pipelined CG
 - depth-one pipeline: p-CG
 - deep pipelines: $p(\ell)$ -CG
- Pipelined BiCGStab
- Pipelined Block-CG/LOBPCG
- Pipelined BiCG/CGLS/LSQR
- Preconditioned pipelined variants
- ► Augmented/deflated/hybrid s-step pipelined methods Yamazaki et al. (2017)

Cornelis et al. (2018)

Ghysels et al. (2013)

- C. & Vanroose (2017)
 - (work in progress)
 - (work in progress)

Ghysels et al. (2014)



Pipelined KSM Numerical stability in finite precision



Pipelined KSM produce identical iterates to classic KSM in exact arithmetic. But finite precision computations introduce roundoff errors, which may lead to:



Pipelined KSM Numerical stability in finite precision



Pipelined KSM produce identical iterates to classic KSM in exact arithmetic. But finite precision computations introduce roundoff errors, which may lead to:

- 1. Delay of convergence: caused by loss of basis orthogonality
- **2. Loss of attainable accuracy:** caused by propagation of local rounding errors in recurrences, and strongly dependent on:
 ▷ pipeline length *l* ▷ system size *N* ▷ polynomial *P_l(A)* (basis *Z_l*)



Rounding errors due to recurrence relations for residual and auxiliary variables:

 $\bar{x}_{i+1} = \bar{x}_i + \bar{\alpha}_i \bar{p}_i + \xi_{i+1}^x, \qquad \bar{r}_{i+1} = \bar{r}_i - \bar{\alpha}_i A \bar{p}_i + \xi_{i+1}^r.$



Rounding errors due to recurrence relations for residual and auxiliary variables:

$$\bar{x}_{i+1} = \bar{x}_i + \bar{\alpha}_i \bar{p}_i + \xi_{i+1}^x, \qquad \bar{r}_{i+1} = \bar{r}_i - \bar{\alpha}_i A \bar{p}_i + \xi_{i+1}^r.$$

Residual deviates from the true residual $b - A\bar{x}_{i+1}$ in finite precision:

$$\begin{aligned} f_{i+1} &= (b - A\bar{x}_{i+1}) - \bar{r}_{i+1} \\ &= b - A(\bar{x}_i + \bar{\alpha}_i \bar{p}_i + \xi_{i+1}^{\times}) - (\bar{r}_i - \bar{\alpha}_i A \bar{p}_i + \xi_{i+1}^{r}) \\ &= f_0 - \sum_{k=0}^{i} (A\xi_{k+1}^{\times} + \xi_{k+1}^{r}). \end{aligned}$$
 Sleijpen & van der Vorst (1995)
 Greenbaum (1997)



Rounding errors due to recurrence relations for residual and auxiliary variables:

$$\bar{x}_{i+1} = \bar{x}_i + \bar{\alpha}_i \bar{p}_i + \xi_{i+1}^x, \qquad \bar{r}_{i+1} = \bar{r}_i - \bar{\alpha}_i A \bar{p}_i + \xi_{i+1}^r.$$

Residual deviates from the true residual $b - A\bar{x}_{i+1}$ in finite precision:

$$\begin{split} f_{i+1} &= (b - A\bar{x}_{i+1}) - \bar{r}_{i+1} \\ &= b - A(\bar{x}_i + \bar{\alpha}_i \bar{p}_i + \xi_{i+1}^x) - (\bar{r}_i - \bar{\alpha}_i A \bar{p}_i + \xi_{i+1}^r) \\ &= f_0 - \sum_{k=0}^i (A\xi_{k+1}^x + \xi_{k+1}^r) . \end{split}$$

Matrix notation: $\mathcal{F}_{i+1} = [f_0, \dots, f_i], \ \Theta_i^x = [0, \xi_1^x, \dots, \xi_{i-1}^x], \ \Theta_i^r = [f_0, \xi_1^r, \dots, \xi_{i-1}^r]$ $\mathcal{F}_{i+1} = -(A\Theta_{i+1}^x + \Theta_{i+1}^r) \ U_{i+1},$

with U_{i+1} an upper triangular matrix with all entries one.



Rounding errors due to recurrence relations for residual and auxiliary variables:

$$\bar{x}_{i+1} = \bar{x}_i + \bar{\alpha}_i \bar{p}_i + \xi_{i+1}^x, \qquad \bar{r}_{i+1} = \bar{r}_i - \bar{\alpha}_i A \bar{p}_i + \xi_{i+1}^r.$$

Residual deviates from the true residual $b - A\bar{x}_{i+1}$ in finite precision:

$$\begin{split} f_{i+1} &= (b - A\bar{x}_{i+1}) - \bar{r}_{i+1} \\ &= b - A(\bar{x}_i + \bar{\alpha}_i \bar{p}_i + \xi_{i+1}^x) - (\bar{r}_i - \bar{\alpha}_i A \bar{p}_i + \xi_{i+1}^r) \\ &= f_0 - \sum_{k=0}^i (A\xi_{k+1}^x + \xi_{k+1}^r) . \end{split}$$

Matrix notation: $\mathcal{F}_{i+1} = [f_0, \dots, f_i], \ \Theta_i^x = [0, \xi_1^x, \dots, \xi_{i-1}^x], \ \Theta_i^r = [f_0, \xi_1^r, \dots, \xi_{i-1}^r]$ $\mathcal{F}_{i+1} = -(A\Theta_{i+1}^x + \Theta_{i+1}^r) \ U_{i+1},$

with U_{i+1} an upper triangular matrix with all entries one.

Only accumulation of local rounding errors in classical CG, no amplification.

Gutknecht & Strakos (2000)I van der Vorst & Ye (2000)

Pipelined KSM Rounding error behavior in pipelined CG

Additional recurrence relations in pipelined CG all introduce local rounding errors:

$$\begin{aligned} \bar{x}_{i+1} &= \bar{x}_i + \bar{\alpha}_i \bar{p}_i + \xi_{i+1}^*, \\ \bar{r}_{i+1} &= \bar{r}_i - \bar{\alpha}_i \bar{s}_i + \xi_{i+1}^r, \\ \bar{p}_i &= \bar{r}_i + \bar{\beta}_i \bar{p}_{i-1} + \xi_i^e, \end{aligned}$$

$$\begin{split} \bar{s}_i &= \bar{w}_i + \bar{\beta}_i \bar{s}_{i-1} + \xi^s_i, \\ \bar{w}_{i+1} &= \bar{w}_i - \bar{\alpha}_i \bar{z}_i + \xi^w_{i+1}, \\ \bar{z}_i &= A \bar{w}_i + \bar{\beta}_i \bar{z}_{i-1} + \xi^z_i, \end{split}$$

Pipelined KSM Rounding error behavior in pipelined CG

Additional recurrence relations in pipelined CG all introduce local rounding errors:

$$\begin{split} \bar{x}_{i+1} &= \bar{x}_i + \bar{\alpha}_i \bar{p}_i + \xi_{i+1}^x, & \bar{s}_i &= \bar{w}_i + \bar{\beta}_i \bar{s}_{i-1} + \xi_i^s, \\ \bar{r}_{i+1} &= \bar{r}_i - \bar{\alpha}_i \bar{s}_i + \xi_{i+1}^r, & \bar{w}_{i+1} &= \bar{w}_i - \bar{\alpha}_i \bar{z}_i + \xi_{i+1}^w, \\ \bar{p}_i &= \bar{r}_i + \bar{\beta}_i \bar{p}_{i-1} + \xi_i^p, & \bar{z}_i &= A \bar{w}_i + \bar{\beta}_i \bar{z}_{i-1} + \xi_i^z, \end{split}$$

Matrix notation: with $f_j = (b - A\bar{x}_j) - \bar{r}_j$, $g_j = A\bar{p}_j - \bar{s}_j$, $h_j = A\bar{r}_j - \bar{w}_j$, $e_j = A\bar{s}_j - \bar{z}_j$

$$\begin{aligned} \mathcal{F}_{j+1} &= (A \Theta_{j+1}^{x} + \Theta_{j+1}^{r}) \, U_{j+1} - \mathcal{G}_{j+1} \mathcal{A}_{j+1} \\ \mathcal{G}_{j+1} &= (A \Theta_{j+1}^{p} + \Theta_{j+1}^{s}) \, \mathcal{B}_{j+1}^{-1} + \mathcal{H}_{j+1} \widetilde{\mathcal{B}_{j+1}^{-1}} \\ \mathcal{H}_{j+1} &= (A \Theta_{j+1}^{u} + \Theta_{j+1}^{w}) \, U_{j+1} - \mathcal{E}_{j+1} \mathcal{A}_{j+1} \\ \mathcal{E}_{j+1} &= (A \Theta_{j+1}^{q} + \Theta_{j+1}^{z}) \, \mathcal{B}_{j+1}^{-1} \end{aligned}$$

Pipelined KSM Rounding error behavior in pipelined CG

Additional recurrence relations in pipelined CG all introduce local rounding errors:

$$\begin{aligned} \bar{\mathbf{x}}_{i+1} &= \bar{\mathbf{x}}_i + \bar{\alpha}_i \bar{\mathbf{p}}_i + \xi_{i+1}^x, & \bar{\mathbf{s}}_i &= \bar{\mathbf{w}}_i + \bar{\beta}_i \bar{\mathbf{s}}_{i-1} + \xi_i^s, \\ \bar{\mathbf{r}}_{i+1} &= \bar{\mathbf{r}}_i - \bar{\alpha}_i \bar{\mathbf{s}}_i + \xi_{i+1}^r, & \bar{\mathbf{w}}_{i+1} &= \bar{\mathbf{w}}_i - \bar{\alpha}_i \bar{\mathbf{z}}_i + \xi_{i+1}^w, \\ \bar{\mathbf{p}}_i &= \bar{\mathbf{r}}_i + \bar{\beta}_i \bar{\mathbf{p}}_{i-1} + \xi_i^p, & \bar{\mathbf{z}}_i &= A \bar{\mathbf{w}}_i + \bar{\beta}_i \bar{\mathbf{z}}_{i-1} + \xi_i^z, \end{aligned}$$

Matrix notation: with $f_j = (b - A\bar{x}_j) - \bar{r}_j$, $g_j = A\bar{p}_j - \bar{s}_j$, $h_j = A\bar{r}_j - \bar{w}_j$, $e_j = A\bar{s}_j - \bar{z}_j$

$$\begin{aligned} \mathcal{F}_{j+1} &= (A\Theta_{j+1}^{x} + \Theta_{j+1}^{r}) U_{j+1} - \mathcal{G}_{j+1} A_{j+1} \\ \mathcal{G}_{j+1} &= (A\Theta_{j+1}^{y} + \Theta_{j+1}^{s}) B_{j+1}^{-1} + \mathcal{H}_{j+1} B_{j+1}^{-1} \\ \mathcal{H}_{j+1} &= (A\Theta_{j+1}^{y} + \Theta_{j+1}^{w}) U_{j+1} - \mathcal{E}_{j+1} A_{j+1} \\ \mathcal{E}_{j+1} &= (A\Theta_{j+1}^{q} + \Theta_{j+1}^{z}) B_{j+1}^{-1} \\ \frac{1}{\beta_{1}} \frac{\beta_{1} \beta_{2}}{\beta_{2}} \cdots \frac{\beta_{1} \beta_{2} \dots \beta_{j}}{\beta_{2} \dots \beta_{j}} \\ \frac{1}{\beta_{1}} \frac{\beta_{1} \beta_{2}}{\beta_{2}} \cdots \frac{\beta_{1} \beta_{2} \dots \beta_{j}}{\beta_{2} \dots \beta_{j}} \\ \frac{1}{\beta_{1}} \frac{\beta_{1} \beta_{2}}{\beta_{2}} \cdots \frac{\beta_{1} \beta_{2} \dots \beta_{j}}{\beta_{2} \dots \beta_{j}} \\ \frac{1}{\beta_{1}} \frac{\beta_{1} \beta_{2}}{\beta_{2}} \cdots \frac{\beta_{1} \beta_{2} \dots \beta_{j}}{\beta_{1} \dots \beta_{j}} \\ \frac{1}{\beta_{1}} \frac{\beta_{1} \beta_{2}}{\beta_{2}} \cdots \frac{\beta_{1} \beta_{2} \dots \beta_{j}}{\beta_{1} \dots \beta_{j}} \\ \frac{1}{\beta_{1}} \frac{\beta_{1} \beta_{2}}{\beta_{2}} \cdots \frac{\beta_{1} \beta_{2} \dots \beta_{j}}{\beta_{1} \dots \beta_{j}} \\ \frac{1}{\beta_{1}} \frac{\beta_{1} \beta_{2}}{\beta_{2}} \cdots \frac{\beta_{1} \beta_{2} \dots \beta_{j}}{\beta_{1} \dots \beta_{j}} \\ \frac{1}{\beta_{1}} \frac{\beta_{1} \beta_{2}}{\beta_{2}} \cdots \frac{\beta_{1} \beta_{2} \dots \beta_{j}}{\beta_{1} \dots \beta_{j}} \\ \frac{1}{\beta_{1}} \frac{\beta_{1} \beta_{2}}{\beta_{1} \dots \beta_{j}} \\ \frac{1}{\beta_{1}} \frac{\beta_{1} \beta_{2}}{\beta_{1} \dots \beta_{j}} \\ \frac{1}{\beta_{1}} \frac{\beta_{1} \beta_{2}}{\beta_{2} \dots \beta_{j}} \\ \frac{1}{\beta_{1}} \frac{\beta_{1} \beta_{2}}{\beta_{1}} \\ \frac{1}{\beta_{1}} \frac{\beta_{1} \beta_{1}}{\beta_{1}} \\ \frac{1}{\beta_{1}} \frac{\beta_{1} \beta_{1}}{\beta$$

$$\beta_i \beta_{i+1} \dots \beta_j = \frac{\|r_j\|^2}{\|r_{i-1}\|^2} \Rightarrow \text{arbitrarily large in CG!}$$

Pipelined KSM Rounding error behavior in pipelined CG

Additional recurrence relations in pipelined CG all introduce local rounding errors:

$$\begin{aligned} \bar{\mathbf{x}}_{i+1} &= \bar{\mathbf{x}}_i + \bar{\alpha}_i \bar{\mathbf{p}}_i + \xi_{i+1}^{\mathsf{x}}, & \bar{\mathbf{s}}_i &= \bar{\mathbf{w}}_i + \bar{\beta}_i \bar{\mathbf{s}}_{i-1} + \xi_i^{\mathsf{s}}, \\ \bar{\mathbf{r}}_{i+1} &= \bar{\mathbf{r}}_i - \bar{\alpha}_i \bar{\mathbf{s}}_i + \xi_{i+1}^{\mathsf{r}}, & \bar{\mathbf{w}}_{i+1} &= \bar{\mathbf{w}}_i - \bar{\alpha}_i \bar{z}_i + \xi_{i+1}^{\mathsf{w}}, \\ \bar{\mathbf{p}}_i &= \bar{\mathbf{r}}_i + \bar{\beta}_i \bar{\mathbf{p}}_{i-1} + \xi_i^{\mathsf{p}}, & \bar{z}_i &= A \bar{\mathbf{w}}_i + \bar{\beta}_i \bar{z}_{i-1} + \xi_i^{\mathsf{z}}, \end{aligned}$$

Matrix notation: with $f_j = (b - A\bar{x}_j) - \bar{r}_j$, $g_j = A\bar{p}_j - \bar{s}_j$, $h_j = A\bar{r}_j - \bar{w}_j$, $e_j = A\bar{s}_j - \bar{z}_j$

$$\begin{aligned} \mathcal{F}_{j+1} &= (\mathcal{A}\Theta_{j+1}^{x} + \Theta_{j+1}^{r}) U_{j+1} - \mathcal{G}_{j+1} \mathcal{A}_{j+1} \\ \mathcal{G}_{j+1} &= (\mathcal{A}\Theta_{j+1}^{p} + \Theta_{j+1}^{s}) \mathcal{B}_{j+1}^{-1} + \mathcal{H}_{j+1} \mathcal{B}_{j+1}^{-1} \\ \mathcal{H}_{j+1} &= (\mathcal{A}\Theta_{j+1}^{u} + \Theta_{j+1}^{w}) U_{j+1} - \mathcal{E}_{j+1} \mathcal{A}_{j+1} \\ \mathcal{E}_{j+1} &= (\mathcal{A}\Theta_{j+1}^{q} + \Theta_{j+1}^{z}) \mathcal{B}_{j+1}^{-1} \\ \mathcal{N} \text{ote:} \\ \beta_{i} \beta_{i+1} \dots \beta_{j} &= \frac{\|l_{\mathcal{I}}^{r}\|^{2}}{\|r_{i-1}\|^{2}} \Rightarrow \text{ arbitrarily large in CG!} \end{aligned} \begin{pmatrix} 0 & \bar{\alpha}_{0} & \bar{\alpha}_{0} & \cdots & \bar{\alpha}_{0} \\ 0 & \bar{\alpha}_{1} & \cdots & \bar{\alpha}_{1} \\ \ddots & \ddots & \vdots \\ 0 & \bar{\alpha}_{j-1} \\ 1 & \bar{\beta}_{1} & \bar{\beta}_{1} \bar{\beta}_{2} & \cdots & \bar{\beta}_{1} \bar{\beta}_{2} \dots \bar{\beta}_{j} \\ 1 & \bar{\beta}_{2} & \bar{\beta}_{2} \dots \bar{\beta}_{j} \\ \vdots & \vdots \\ 1 & \bar{\beta}_{j} & 1 & \bar{\beta}_{j} \\ \vdots & \vdots \\ 1 & 0 & 1 & 1 \\ \end{pmatrix}$$

Amplification of local rounding errors possible, depending on $\bar{\alpha}_i$'s and $\bar{\beta}_i$'s.

Carson et al. (2017) C. & Vanroose (2017)



Numerical stability analysis of p(/)-CG in a nutshell



Numerical stability analysis of p(/)-CG in a nutshell

U True basis vector $\bar{\mathbf{v}}_{j+1}$ satisfies:

$$Aar{V}_{j} = ar{V}_{j+1}ar{H}_{j+1,j} + (ar{V}_{j+1} - ar{V}_{j+1})ar{\Delta}_{j+1,j}$$



Numerical stability analysis of p(/)-CG in a nutshell

U True basis vector $\bar{\mathbf{v}}_{j+1}$ satisfies:

$$Aar{V}_j=ar{V}_{j+1}ar{H}_{j+1,j}+(ar{V}_{j+1}-ar{V}_{j+1})ar{\Delta}_{j+1,j}$$

2 Computed basis vector \bar{v}_{j+1} satisfies:

 $ar{Z}_j = ar{V}_j ar{G}_j + \Theta_j^v$



Numerical stability analysis of p(/)-CG in a nutshell

True basis vector $\bar{\mathbf{v}}_{j+1}$ satisfies: $A\bar{V}_j = \bar{V}_{j+1}\bar{H}_{j+1,j} + (\bar{\mathbf{v}}_{j+1} - \bar{V}_{j+1})\bar{\Delta}_{j+1,j}$ Computed basis vector \bar{v}_{j+1} satisfies: $\bar{Z}_j = \bar{V}_j\bar{G}_j + \Theta_j^{v}$ Computed basis vector \bar{z}_{j+1} satisfies: $A\bar{Z}_i = \bar{Z}_{i+1}\bar{B}_{i+1,i} + \Theta_i^{z}$



Numerical stability analysis of p(/)-CG in a nutshell

 $\stackrel{\textbf{l}}{=}$ True basis vector $\overline{\mathbf{v}}_{j+1}$ satisfies: $A\bar{V}_{i} = \bar{V}_{i+1}\bar{H}_{i+1,i} + (\bar{V}_{i+1} - \bar{V}_{i+1})\bar{\Delta}_{i+1,i}$ Computed basis vector \bar{v}_{j+1} satisfies: $\bar{Z}_i = \bar{V}_i \bar{G}_i + \Theta_i^v$ Computed basis vector \bar{z}_{j+1} satisfies: $A\overline{Z}_i = \overline{Z}_{i+1}\overline{B}_{i+1}i + \Theta_i^z$ Using 1, 2 and 3 the gap $\mathcal{F}_{j+1} = m{V}_{j+1} - m{V}_{j+1}$ can be calculated as $\mathcal{F}_{i+1} = (\Theta_i^{\bar{z}} \bar{G}_i^{-1} - A \Theta_i^{\bar{v}} \bar{G}_i^{-1} + \Theta_{i+1}^{\bar{v}} \bar{B}_{i+1,i} \bar{G}_i^{-1}) \bar{\Delta}_{i+1,i}^{-1}$



Numerical stability analysis of p(/)-CG in a nutshell



Pipelined KSM Rounding error behavior in p(I)-CG



 Maximum norm ||G_i⁻¹||_{max} provides a measure for the impact of local rounding error propagation on maximal attainable accuracy in p(1)-CG

Pipelined KSM Rounding error behavior in p(I)-CG



- Maximum norm ||G_i⁻¹||_{max} provides a measure for the impact of local rounding error propagation on maximal attainable accuracy in p(1)-CG
- Rounding error analysis explains loss of attainable accuracy w.r.t.
 ▷ pipeline length *I* ▷ system size *N* ▷ polynomial *P_I(A)* (basis *Z_i*)

Pipelined KSM Residual replacement in pipelined CG

• Core idea: improve accuracy by replacing \bar{r}_i , \bar{s}_i , \bar{w}_i and \bar{z}_i by their true values in selected iterations (causing minor computational overhead):

 $\bar{r}_i = \mathrm{fl}(b - A\bar{x}_i), \quad \bar{w}_i = \mathrm{fl}(A\bar{r}_i), \quad \bar{s}_i = \mathrm{fl}(A\bar{p}_i), \quad \bar{z}_i = \mathrm{fl}(A\bar{s}_i).$

van der Vorst & Ye (2000)

Pipelined KSM Residual replacement in pipelined CG

• Core idea: improve accuracy by replacing \bar{r}_i , \bar{s}_i , \bar{w}_i and \bar{z}_i by their true values in selected iterations (causing minor computational overhead):

 $\bar{r}_i = \mathrm{fl}(b - A\bar{x}_i), \quad \bar{w}_i = \mathrm{fl}(A\bar{r}_i), \quad \bar{s}_i = \mathrm{fl}(A\bar{p}_i), \quad \bar{z}_i = \mathrm{fl}(A\bar{s}_i).$

🖻 van der Vorst & Ye (2000)

 Replacement criterion: choose when to replace based on (an estimate of) the gap ||f_i|| = ||(b − Ax̄_i) − r̄_i||:

 $\|f_{i-1}\| \leq \tau \|\overline{r}_{i-1}\|$ and $\|f_i\| > \tau \|\overline{r}_i\|$ with $\tau = \sqrt{\epsilon}$.

- Replace sufficiently often such that $||f_i||$ remains small
- ▶ Don't replace when $\|\bar{r}_i\|$ is too small, which may cause delay of convergence

Sleijpen & van der Vorst (1996) Strakos & Tichy (2002)

Pipelined KSM Residual replacement in pipelined CG

• Core idea: improve accuracy by replacing \bar{r}_i , \bar{s}_i , \bar{w}_i and \bar{z}_i by their true values in selected iterations (causing minor computational overhead):

 $\bar{r}_i = \mathrm{fl}(b - A\bar{x}_i), \quad \bar{w}_i = \mathrm{fl}(A\bar{r}_i), \quad \bar{s}_i = \mathrm{fl}(A\bar{p}_i), \quad \bar{z}_i = \mathrm{fl}(A\bar{s}_i).$

🖻 van der Vorst & Ye (2000)

 Replacement criterion: choose when to replace based on (an estimate of) the gap ||f_i|| = ||(b − Ax̄_i) − r̄_i||:

 $\|f_{i-1}\| \leq \tau \|\overline{r}_{i-1}\|$ and $\|f_i\| > \tau \|\overline{r}_i\|$ with $\tau = \sqrt{\epsilon}$.

- Replace sufficiently often such that $||f_i||$ remains small
- ▶ Don't replace when $\|\bar{r}_i\|$ is too small, which may cause delay of convergence

Sleijpen & van der Vorst (1996) Strakos & Tichy (2002)

Automated residual replacements: estimate of || *f_i* || is computed at runtime (inexpensive). In many cases accuracy is significantly improved.

 ■ C. & Vanroose (2017)



Pipelined KSM Residual replacement in pipelined CG

- ▶ PETSc implementation using MPICH-3.1.3 communication
- Benchmark problem: 2D Laplacian model, 1,000,000 unknowns
- ► System specs: 20 nodes, two 6-core Intel Xeon X5660 Nehalem 2.8GHz CPUs/node



Speedup over single-node CG





Pipelined KSM Numerically stable variant of p(*I*)-CG

Use idea similar to residual replacement to improve stability: replace recurrence for v_{i+1} by the Arnoldi relation:

Original p(1)-CG:

(+

$$v_{j+1} = \left(z_{j+1} - \sum_{k=j-2l+1}^{j} g_{k,j+1} v_k\right) / g_{j+1,j+1}$$



Pipelined KSM Numerically stable variant of p(/)-CG

Use idea similar to residual replacement to improve stability: replace recurrence for v_{i+1} by the Arnoldi relation:

Original p(1)-CG:

1+

$$v_{j+1} = \left(z_{j+1} - \sum_{k=j-2l+1}^{j} g_{k,j+1} v_k\right) / g_{j+1,j+1}$$

Stabilized p(I)-CG:

 $\mathbf{v}_{j+1} = (\mathbf{A}\mathbf{v}_j - \gamma_j\mathbf{v}_j - \delta_{j-1}\mathbf{v}_{j-1})/\delta_j$



Pipelined KSM Numerically stable variant of p(/)-CG

Use idea similar to residual replacement to improve stability: replace recurrence for v_{i+1} by the Arnoldi relation:

Original p(1)-CG:

1+

$$v_{j+1} = \left(z_{j+1} - \sum_{k=j-2l+1}^{j} g_{k,j+1} v_k\right) / g_{j+1,j+1}$$

Stabilized p(I)-CG:

 $\mathbf{v}_{j+1} = (\mathbf{A}\mathbf{v}_j - \gamma_j\mathbf{v}_j - \delta_{j-1}\mathbf{v}_{j-1})/\delta_j$

Attainable accuracy is improved for all pipeline lengths I, but extra SpMV increases computation time

⇒ trade-off between numerical stability and performance





Conclusions and takeaways Wrap-up of this talk

- Pipelined Krylov subspace methods are a promising approach to reduce synchronization cost in linear solvers for large-scale problems
 - By adding *auxiliary variables* and recurrences it is possible to *hide* communication latency behind computational kernels
 - Deep pipelines allow to hide global reductions behind multiple SpMV's
 - Asynchronous implementation: dot-products can take multiple iterations to complete, in an overlapping manner
 - Improved scaling over classic KSMs in strong scaling limit, where global reduction latencies rise and volume of computations per core diminishes



Conclusions and takeaways Wrap-up of this talk

- Pipelined Krylov subspace methods are a promising approach to reduce synchronization cost in linear solvers for large-scale problems
 - By adding *auxiliary variables* and recurrences it is possible to *hide* communication latency behind computational kernels
 - Deep pipelines allow to hide global reductions behind multiple SpMV's
 - Asynchronous implementation: dot-products can take multiple iterations to complete, in an overlapping manner
 - Improved scaling over classic KSMs in strong scaling limit, where global reduction latencies rise and volume of computations per core diminishes
- Finite precision behavior of communication reducing and hiding algorithms should be carefully monitored!
 - ► Rounding error analysis allows to explain observed loss of attainable accuracy
 - Residual replacement -type techniques can be applied to improve numerical stability, but at a (slight) increase in computational cost



Conclusions and takeaways Related work

- Work in progress (not covered in this talk):
 - Impact of hard faults & soft-errors on pipelined Krylov subspace methods (INRIA Bordeaux)
 - Resilience of pipelined Krylov subspace methods to system noise (UChicago/ETHZ/Rice/UIIInois)
 - Pipelined Lanczos for eigenvalue calculation in graph partitioning algorithms (UAntwerp/LBNL)
 - Pipelined Blocked Krylov subspace methods for systems with multiple rhs (UAntwerp/INRIA Bordeaux)
 - Communication avoiding/hiding LOBPCG for eigenvalue calculations (UAntwerp/LBNL)
- Many interesting open problems and challenges remain as we push toward exascale-level computing!



Conclusions and takeaways Contributions to PETSc

Open source HPC linear algebra toolkit: https://www.mcs.anl.gov/petsc/

T	C petsc	petsc / PETSc / petsc		
۹	Cverview	Overview		
+	Source	▲ SSH v git@bitbucket.orgpetsc/petsc.git		
	Commits	Last undated an hour ann	14	107
	🕼 Branches	Website http://mcsanlgov/petsc	Open PRs	Watchers
	N Pull requests	Language C Access level Read	99+	135
	Ø Pipelines		Branches	Forks

- ► KSPPGMRES: pipelined GMRES (thanks to J. Brown)
- ► KSPPIPECG: pipelined CG
- ► KSPPIPECGRR: pipelined CG with automated residual replacement
- ► KSPPIPELCG: pipelined CG with deep pipelines
- KPPPIPECR: pipelined conjugate residuals
- ► KSPGROPPCG: asynchronous CG variant by W. Gropp and collaborators
- KSPPIPEBCGS: pipelined BiCGStab



Conclusions and takeaways Contributions to PETSc

Open source HPC linear algebra toolkit: https://www.mcs.anl.gov/petsc/





Thank you!

siegfried.cools@uantwerp.be

https://www.uantwerpen.be/en/staff/siegfried-cools

Publications



- P. Ghysels, T. J. Ashby, K. Meerbergen, and W. Vanroose, *Hiding Global Communication Latency in the GMRES Algorithm on Massively Parallel Machines* SIAM J. Sci. Comput., 35(1), pp. C48C71, 2013.
- P. Ghysels and W. Vanroose, *Hiding global synchronization latency in the preconditioned Conjugate Gradient algorithm*, Parallel Computing, 40(7), pp. 224-238, 2014.
- S. Cools, E.F. Yetkin, E. Agullo, L. Giraud, W. Vanroose, Analyzing the effect of local rounding error propagation on the maximal attainable accuracy of the pipelined Conjugate Gradient method. SIAM J. on Matrix Anal. Appl., 39(1), pp. 426-450, 2018.



S. Cools, W. Vanroose, The communication-hiding pipelined BiCGStab method for the parallel solution of large unsymmetric linear systems. Parallel Computing, 65, pp. 1-20, Elsevier, 2017.



J. Cornelis, S. Cools, W. Vanroose, *The communication-hiding Conjugate Gradient method with deep pipelines.* Submitted to SIAM J. Sci. Comput., 2018, Preprint available at arXiv:1801.04728.



S. Cools, Numerical stability analysis of the class of communication hiding pipelined Conjugate Gradient methods. Submitted to SIAM J. Sci. Comput., 2018, Preprint available at arXiv:1804.02962.