Analyzing and Improving Attainable Accuracy for the Communication Hiding Pipelined Conjugate Gradient Method

SIAM Conference on Computational Science and Engineering (CSE'19)

Spokane Convention Center, WA, US, February 25 - March 1, 2019

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

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

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

Universiteit Antwerpen

Communication is hard for humans

Introducing some buzzwords you will frequently hear during this talk ...

"*synchronization bottleneck*" "*communication latency*" / "*error propagation*"

Communication is hard for computers

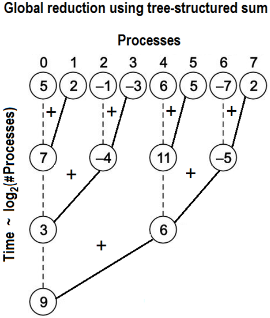

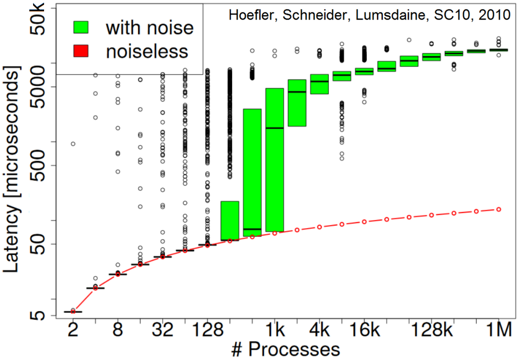

Data movement (communication) is much more time consuming than flops (computations), so reducing time spent communicating data is essential for HPC

$\Rightarrow$ Communication avoiding / Communication hiding

Communication reducing methods

High communication cost has motivated several approaches to reducing the global synchronization bottleneck in Krylov subspace methods:

Avoiding communication: s-step Krylov subspace methods *

[A. Chronopoulous, E. de Sturler, J. Demmel, M. Hoemmen, E. Carson, L. Grigori, J. Erhel, ...]

- Compute iterations in blocks of s, allows use of matrix power kernels
- 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, P. Luszczek, ...]

- Introduce auxiliary (basis) vectors to decouple SpMV and inner products
- Enables overlapping of communication and computations

* All methods are equivalent to their corresponding Krylov subspace methods in exact arithmetic

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

$$x_i \in x_0 + \mathcal{K}_i(A, r_0) = x_0 + \text{span}\{r_0, Ar_0, A^2 r_0, \dots, A^{i-1} r_0\}, \quad r_i = b - Ax_i.$$

- minimize certain error measure over Krylov subspace $\mathcal{K}_i(A, r_0)$
- Krylov subspace methods:

  Conjugate Gradients (CG), Lanczos, GMRES, MinRES, BiCG, 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
     - $\mathcal{O}(\text{N})$ flops
     - global synchronization (MPI_Allreduce)
  ii. SpMV
     - $\mathcal{O}(\text{nnz})$ flops
     - neighbor communication only
  iii. axpy
     - $\mathcal{O}(\text{N})$ flops
     - no communication

Classic Conjugate Gradients (CG)

Algorithm CG

procedure CG(A, b, x_0)
  $r_0 := b - Ax_0$; $p_0 = r_0$
  for $i = 0, \dots$ **do**
    $s_i := Ap_i$
    $\alpha_i := (r_i, r_i) \,/\, (s_i, p_i)$
    $x_{i+1} := x_i + \alpha_i p_i$
    $r_{i+1} := r_i - \alpha_i s_i$
    $\beta_{i+1} := (r_{i+1}, r_{i+1}) \,/\, (r_i, r_i)$
    $p_{i+1} := r_{i+1} + \beta_{i+1} p_i$
  end for
end procedure

dot-pr
SpMV
axpy

Hestenes & Stiefel (1952)

i. dot-products
- 2 global reductions: latency dominated
- time scales as $\log_2(\#\textit{partitions})$

ii. SpMV
- computationally expensive
- good scaling (minor communication)

iii. axpy's
- vector operations (recurrences)
- perfect scaling (no communication)

Essentially sequential operations (line-per-line execution)

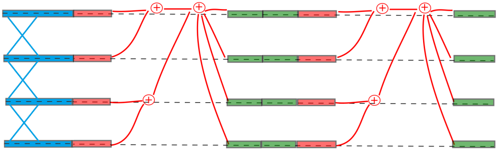

Pipelined Conjugate Gradients

Algorithm Pipelined CG

```
procedure PIPE-CG(A, b, x_0)
    r_0 := b − A x_0; w_0 := A r_0
    for i = 0, . . . do
        γ_i := (r_i, r_i)
        δ := (w_i, r_i)
        q_i := A w_i
        if i > 0 then
            β_i := γ_i/γ_{i−1}; α_i := (δ/γ_i − β_i/α_{i−1})^{−1}
        else
            β_i := 0; α := γ_i/δ
        end if
        z_i := q_i + β_i z_{i−1}
        s_i := w_i + β_i s_{i−1}
        p_i := r_i + β_i p_{i−1}
        x_{i+1} := x_i + α_i p_i
        r_{i+1} := r_i − α_i s_i
        w_{i+1} := w_i − α_i z_i
    end for
end procedure
```

dot-pr
SpMV
axpy

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 (axpy's) for the auxiliary vectors $s_i = Ap_i$, $w_i = Ar_i$, $z_i = As_i$

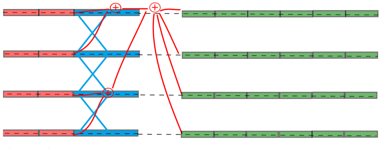

Deep ℓ-length pipelined CG

Classic KSM:

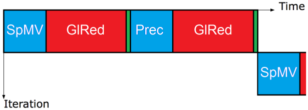

Pipelined KSM:

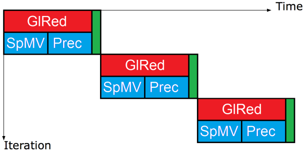

Deep pipelined KSM:

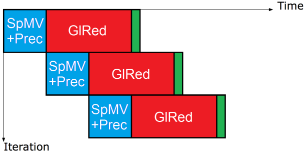

Pipelined "D-Lanczos" — Saad (2003)

Consider the Lanczos relation

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

with A symmetric, $V_{i+1} = [v_0, v_1, \ldots, v_i]$ the Krylov subspace basis and $T_{i+1,i}$ a symmetric tridiagonal matrix

$$T_{i+1,i} = \begin{pmatrix} \gamma_0 & \delta_0 & & \\ \delta_0 & \gamma_1 & \ddots & \\ & \ddots & \ddots & \delta_{i-2} \\ & & \delta_{i-2} & \gamma_{i-1} \\ & & & \delta_{i-1} \end{pmatrix}.$$

Deep ℓ-length pipelined CG

Classic KSM:

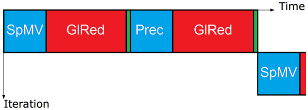

Pipelined KSM:

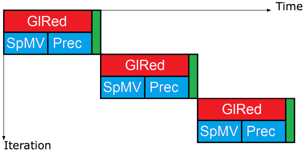

Deep pipelined KSM:

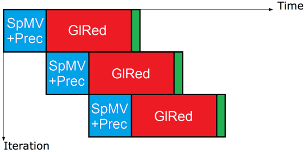

Pipelined "D-Lanczos" Saad (2003)

Introduce the auxiliary Krylov subspace basis $Z_{i+1} = [z_0, z_1, \ldots, z_i]$ that runs l SpMVs ahead of the basis V_{i-l+1} as

$$z_i := \begin{cases} v_0 & j = 0, \\ P_i(A)v_0 & 0 < i \leq l, \\ P_l(A)v_{i-l} & i > l, \end{cases}$$

with polynomials $P_l(t)$ of fixed order l

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

where l is the pipeline length.

Ghysels et al. (2013)

Deep ℓ-length pipelined CG

- Applying $P_l(A)$ to $AV_i = V_{i+1} T_{i+1,i}$ yields a Lanczos-type relation

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

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

- Auxiliary basis vectors are computed using a three-term recurrence relation

$$z_{i+1} = (\underbrace{Az_i}_{\text{SpMV}} - \gamma_{i-l} z_i - \delta_{i-l-1} z_{i-1})/\delta_{i-l}$$

- Basis transformation. Z_i and V_i both span i-th Krylov subspace, thus $\exists$ an upper triangular basis transformation matrix G_i with

$$Z_i = V_i G_i.$$

- Band structure of G_i. Matrix G_i has only $2l+1$ nonzero diagonals

$$\begin{aligned} g_{j,i} &= (z_i, v_j) = (P_l(A) v_{i-l}, v_j) \\ &= (v_{i-l}, P_l(A) v_j) = g_{i-l,j+l}. \end{aligned}$$

- Original basis vectors are computed using a multi-term recurrence relation

$$v_{i-l+1} = \left(z_{i-l+1} - \sum_{j=i-3l+1}^{i-l} g_{j,i-l+1} \, v_j \right) / g_{i-l+1,i-l+1}.$$

Cornelis et al. (2018)

Deep ℓ-length pipelined CG

Algorithm 1 l-length pipelined p(l)-CG **Input:** A, b, x_0, l, m, τ

$r_0 := b - Ax_0$; $v_0 := r_0/\|r_0\|_2$; $z_0 := v_0$; $g_{0,0} := 1$;
for $i = 0, \dots, m+l$ **do**
  $z_{i+1} := \begin{cases} (A - \sigma_i I) z_i, & i < l \\ A z_i, & i \geq l \end{cases}$
  if $i \geq l$ **then**
    $g_{j,i-l+1} := (g_{j,i-l+1} - \sum_{k=i-3l+1}^{j-1} g_{k,j} g_{k,i-l+1})/g_{j,j}$; $j = i-2l+2, \dots, i-l$
    $g_{i-l+1,i-l+1} := \sqrt{g_{i-l+1,i-l+1} - \sum_{k=i-3l+1}^{i-l} g_{k,i-l+1}^2}$;
    # Check for breakdown and restart if required
    if $i < 2l$ **then**
      $\gamma_{i-l} := (g_{i-l,i-l+1} + \sigma_{i-l} g_{i-l,i-l} - g_{i-l-1,i-l}\delta_{i-l-1})/g_{i-l,i-l}$;
      $\delta_{i-l} := g_{i-l+1,i-l+1}/g_{i-l,i-l}$;
    else
      $\gamma_{i-l} := (g_{i-l,i-l}\gamma_{i-2l} + g_{i-l,i-l+1}\delta_{i-2l} - g_{i-l-1,i-l}\delta_{i-l-1})/g_{i-l,i-l}$;
      $\delta_{i-l} := (g_{i-l+1,i-l+1}\delta_{i-2l})/g_{i-l,i-l}$;
    end if
    $v_{i-l+1} := (z_{i-l+1} - \sum_{j=i-3l+1}^{i-l} g_{j,i-l+1} v_j)/g_{i-l+1,i-l+1}$;
    $z_{i+1} := (z_{i+1} - \gamma_{i-l} z_i - \delta_{i-l-1} z_{i-1})/\delta_{i-l}$;
  end if
  $g_{j,i+1} := \begin{cases} (z_{i+1}, v_j); & j = \max(0, i-2l+1), \dots, i-l+1 \\ (z_{i+1}, z_j); & j = i-l+2, \dots, i+1 \end{cases}$
  if $i = l$ **then**
    $\eta_0 := \gamma_0$; $\zeta_0 := \|r_0\|_2$; $p_0 := v_0/\eta_0$;
  else if $i \geq l+1$ **then**
    $\lambda_{i-l} := \delta_{i-l-1}/\eta_{i-l-1}$;
    $\eta_{i-l} := \gamma_{i-l} - \lambda_{i-l}\delta_{i-l-1}$;
    $\zeta_{i-l} = -\lambda_{i-l}\zeta_{i-l-1}$;
    $p_{i-l} = (v_{i-l} - \delta_{i-l-1} p_{i-l-1})/\eta_{i-l}$;
    $x_{i-l} = x_{i-l-1} + \zeta_{i-l-1} p_{i-l-1}$;
    if $|\zeta_{i-l}|/\|r_0\| < \tau$ **then** RETURN; **end if**
  end if
end for

Deep ℓ-length pipelined CG

Algorithm 1 l-length pipelined p(l)-CG **Input:** A, b, x_0, l, m, τ

$r_0 := b - Ax_0$; $v_0 := r_0/\|r_0\|_2$; $z_0 := v_0$; $g_{0,0} := 1$;
for $i = 0, \ldots, m+l$ **do**
  $z_{i+1} := \begin{cases} (A - \sigma_i I) z_i, & i < l \\ A z_i, & i \geq l \end{cases}$
  if $i \geq l$ **then**
    $g_{j,i-l+1} := (g_{j,i-l+1} - \sum_{k=i-3l+1}^{j-1} g_{k,j} g_{k,i-l+1})/g_{j,j}$; $j = i - 2l + 2, \ldots, i$
    $g_{i-l+1,i-l+1} := \sqrt{g_{i-l+1,i-l+1} - \sum_{k=i-3l+1}^{i-l} g_{k,i-l+1}^2}$;
    # Check for breakdown and restart if required
    if $i < 2l$ **then**
      $\gamma_{i-l} := (g_{i-l,i-l+1} + \sigma_{i-l} g_{i-l,i-l} - g_{i-l-1,i-l}\delta_{i-l-1})/g_{i-l,i-l}$;
      $\delta_{i-l} := g_{i-l+1,i-l+1}/g_{i-l,i-l}$;
    else
      $\gamma_{i-l} := (g_{i-l,i-l}\gamma_{i-2l} + g_{i-l,i-l+1}\delta_{i-2l} - g_{i-l-1,i-l}\delta_{i-l-1})/g_{i-l,i-l}$;
      $\delta_{i-l} := (g_{i-l+1,i-l+1}\delta_{i-2l})/g_{i-l,i-l}$;
    end if
    $v_{i-l+1} := (z_{i-l+1} - \sum_{j=i-3l+1}^{i-l} g_{j,i-l+1} v_j)/g_{i-l+1,i-l+1}$;
    $z_{i+1} := (z_{i+1} - \gamma_{i-l} z_i - \delta_{i-l-1} z_{i-1})/\delta_{i-l}$;
  end if
  $g_{j,i+1} := \begin{cases} (z_{i+1}, v_j); & j = \max(0, i-2l+1), \ldots, i-l+1 \\ (z_{i+1}, z_j); & j = i-l+2, \ldots, i+1 \end{cases}$
  if $i = l$ **then**
    $\eta_0 := \gamma_0$; $\zeta_0 := \|r_0\|_2$; $p_0 := v_0/\eta_0$;
  else if $i \geq l+1$ **then**
    $\lambda_{i-l} := \delta_{i-l-1}/\eta_{i-l-1}$;
    $\eta_{i-l} := \gamma_{i-l} - \lambda_{i-l}\delta_{i-l-1}$;
    $\zeta_{i-l} = -\lambda_{i-l}\zeta_{i-l-1}$;
    $p_{i-l} = (v_{i-l} - \delta_{i-l-1} p_{i-l-1})/\eta_{i-l}$;
    $x_{i-l} = x_{i-l-1} + \zeta_{i-l-1} p_{i-l-1}$;
    if $|\zeta_{i-l}|/\|r_0\| < \tau$ **then** RETURN; **end if**
  end if
end for

SpMV (+ preconditioner)

- 1 SpMV on z_i per iteration

Recurrence relations for V_{i-l+1} and Z_{i+1} basis vectors

- computation: $2l + 2$ axpy's
- storage: $3l + 2$ basis vectors

$$z_0 \;\; z_1 \;\; \ldots \;\; z_{i-l} \;\; \overbrace{z_{i-l+1} \;\; \ldots \;\; z_{i-1} \;\; z_i \;\; z_{i+1}}^{l+1 \text{ vectors}}$$
$$\underbrace{v_0 \;\; v_1 \;\; \ldots \;\; v_{i-3l+1} \;\; \ldots \;\; v_{i-l} \;\; v_{i-l+1}}_{2l+1 \text{ vectors}}$$

Dot-products

- $2l + 1$ band structure of G_i
- one global reduction phase is initiated per iteration

Deep ℓ-length pipelined CG

Algorithm 1 l-length pipelined p(l)-CG **Input:** A, b, x_0, l, m, τ

$r_0 := b - Ax_0$; $v_0 := r_0/\|r_0\|_2$; $z_0 := v_0$; $g_{0,0} := 1$;
for $i = 0, \dots, m+l$ **do**
  $z_{i+1} := \begin{cases} (A - \sigma_i I) z_i, & i < l \\ A z_i, & i \geq l \end{cases}$
  if $i \geq l$ **then**
    $g_{j,i-l+1} := (g_{j,i-l+1} - \sum_{k=i-3l+1}^{j-1} g_{k,j} g_{k,i-l+1})/g_{j,j}$; $j = i-2l+2, \dots, i$
    $g_{i-l+1,i-l+1} := \sqrt{g_{i-l+1,i-l+1} - \sum_{k=i-3l+1}^{i-l} g_{k,i-l+1}^2}$;
    # Check for breakdown and restart if required
    if $i < 2l$ **then**
      $\gamma_{i-l} := (g_{i-l,i-l+1} + \sigma_{i-l} g_{i-l,i-l} - g_{i-l-1,i-l}\delta_{i-l-1})/g_{i-l,i-l}$;
      $\delta_{i-l} := g_{i-l+1,i-l+1}/$
    else
      $\gamma_{i-l} := (g_{i-l,i-l}\gamma_{i-2}$
      $\delta_{i-l} := (g_{i-l+1,i-l+1}$
    end if
    $v_{i-l+1} := (z_{i-l+1} - \sum_{j}$
    $z_{i+1} := (z_{i+1} - \gamma_{i-l} z_i -$
  end if
  $g_{j,i+1} := \begin{cases} (z_{i+1}, v_j); \\ (z_{i+1}, z_j); \end{cases}$ $j =$
  if $i = l$ **then**
    $\eta_0 := \gamma_0$; $\zeta_0 := \|r_0\|_2$; $p_0 := v_0/\eta_0$;
  else if $i \geq l+1$ **then**
    $\lambda_{i-l} := \delta_{i-l-1}/\eta_{i-l-1}$;
    $\eta_{i-l} := \gamma_{i-l} - \lambda_{i-l}\delta_{i-l-1}$;
    $\zeta_{i-l} = -\lambda_{i-l}\zeta_{i-l-1}$;
    $p_{i-l} = (v_{i-l} - \delta_{i-l-1} p_{i-l-1})/\eta_{i-l}$;
    $x_{i-l} = x_{i-l-1} + \zeta_{i-l-1} p_{i-l-1}$;
    if $|\zeta_{i-l}|/\|r_0\| < \tau$ **then** RETURN; **end if**
  end if
end for

Results of global sync. are needed l iterations later to update $G_{:,i-l+1}$

Each global reduction is overlapped by ℓ SpMVs

SpMV (+ preconditioner)

- 1 SpMV on z_i per iteration

Recurrence relations for V_{i-l+1} and Z_{i+1} basis vectors

- computation: $2l + 2$ axpy's
- storage: $3l + 2$ basis vectors

$$z_0 \; z_1 \; \dots \; z_{i-l} \; \overbrace{z_{i-l+1} \; \dots \; z_{i-1} \; z_i \; z_{i+1}}^{l+1 \text{ vectors}}$$
$$v_0 \; v_1 \; \dots \; \underbrace{v_{i-3l+1} \; \dots \; v_{i-l} \; v_{i-l+1}}_{2l+1 \text{ vectors}}$$

Dot-products

- $2l + 1$ band structure of G_i
- one global reduction phase is initiated per iteration

Parallel performance of pipelined CG

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

Per node: Two 6-core Intel Xeon X5660 Nehalem 2.80 GHz - 2D Poisson (5pt) - 1 million unknowns

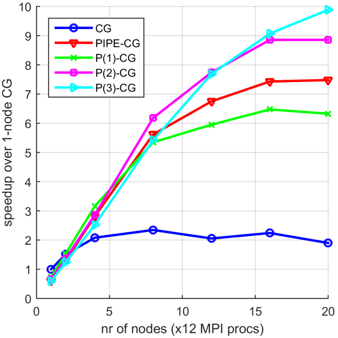

Per node: Two 14-core Intel E5-2680v4 Broadwell 2.40 GHz - 2D Poisson (5pt) - 3 million unknowns

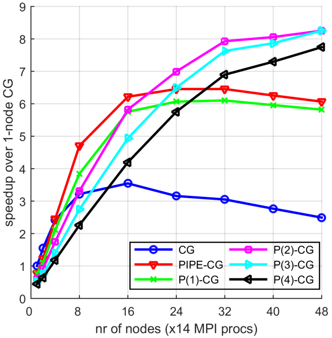

Cornelis et al. (2018)

Numerical stability in finite precision

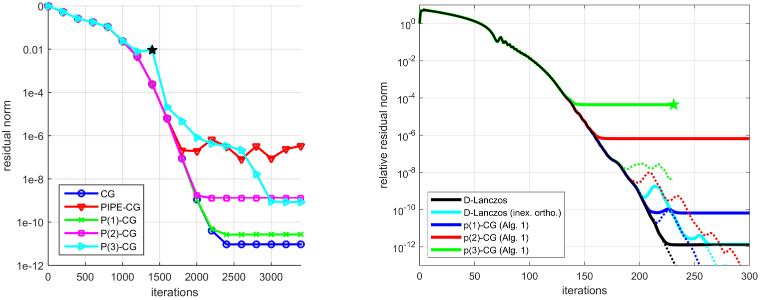

Pipelined CG produces identical iterates to classic CG in exact arithmetic; but ...

Finite precision computations introduce roundoff errors that may lead to

1. *Delayed convergence* due to loss of basis orthogonality
2. *Loss of attainable accuracy* due to propagation of local rounding errors introduced by the recurrence relations

Analyzing rounding error behavior in CG

Rounding errors due to recurrence relations for residual and solution update:

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

Computed residual $\bar{r}_i$ deviates from the true residual $b - A\bar{x}_i$ in finite precision:

$$\begin{aligned} (b - A\bar{x}_{i+1}) - \bar{r}_{i+1} &= b - A(\bar{x}_i + \bar{\alpha}_i \bar{p}_i + \xi^x_{i+1}) - (\bar{r}_i - \bar{\alpha}_i A\bar{p}_i + \xi^r_{i+1}) \\ &= \sum_{k=0}^{i+1} (A\xi^x_k + \xi^r_k)\,. \end{aligned}$$

Sleijpen et al. (1995)
Greenbaum (1997)

Matrix notation: $\bar{R}_{i+1} = [\bar{r}_0, \ldots, \bar{r}_i]$, $\bar{X}_{i+1} = [\bar{x}_0, \ldots, \bar{x}_i]$, Θ^x_i, Θ^r_i rounding errors

$$(B - A\bar{X}_{i+1}) - \bar{R}_{i+1} = (A\Theta^x_{i+1} + \Theta^r_{i+1})\, E_{i+1},$$

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

Accumulation of local rounding errors in classic CG, but no amplification.

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

Analyzing pipelined CG by Ghysels et al.

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

$$
\begin{aligned}
\bar{x}_{i+1} &= \bar{x}_i + \bar{\alpha}_i \bar{p}_i + \xi^x_{i+1}, & \bar{s}_i &= \bar{w}_i + \bar{\beta}_i \bar{s}_{i-1} + \xi^s_i, \\
\bar{r}_{i+1} &= \bar{r}_i - \bar{\alpha}_i \bar{s}_i + \xi^r_{i+1}, & \bar{w}_{i+1} &= \bar{w}_i - \bar{\alpha}_i \bar{z}_i + \xi^w_{i+1}, \\
\bar{p}_i &= \bar{r}_i + \bar{\beta}_i \bar{p}_{i-1} + \xi^p_i, & \bar{z}_i &= A\bar{w}_i + \bar{\beta}_i \bar{z}_{i-1} + \xi^z_i,
\end{aligned}
$$

The gap on the residual is coupled to the gaps on the auxiliary variables:

$$
\begin{aligned}
(B - A\bar{X}_i) - \bar{R}_i = {} & (A\Theta^{\bar{x}}_i + \Theta^{\bar{r}}_i)\, E_i + (A\Theta^{\bar{p}}_i + \Theta^{\bar{s}}_i)\, \bar{\mathcal{B}}_i^{-1} \bar{\mathcal{A}}_i \\
& + (A\Theta^{\bar{u}}_i + \Theta^{\bar{w}}_i)\, E_i \bar{\mathcal{B}}_i^{-1} \bar{\mathcal{A}}_i + (A\Theta^{\bar{q}}_i + \Theta^{\bar{z}}_i)\, \bar{\mathcal{B}}_i^{-1} \bar{\mathcal{A}}_i \bar{\mathcal{B}}_i^{-1} \bar{\mathcal{A}}_i
\end{aligned}
$$

with $\bar{\mathcal{A}}_i = \begin{pmatrix} 0 & \bar{\alpha}_0 & \bar{\alpha}_0 & \cdots & \bar{\alpha}_0 \\ & 0 & \bar{\alpha}_1 & \cdots & \bar{\alpha}_1 \\ & & \ddots & & \vdots \\ & & & 0 & \bar{\alpha}_{i-2} \\ & & & & 0 \end{pmatrix}$, $\bar{\mathcal{B}}_i^{-1} = \begin{pmatrix} 1 & \bar{\beta}_1 & \bar{\beta}_1\bar{\beta}_2 & \cdots & \bar{\beta}_1\bar{\beta}_2\ldots\bar{\beta}_{i-1} \\ & 1 & \bar{\beta}_2 & & \bar{\beta}_2\ldots\bar{\beta}_{i-1} \\ & & \ddots & & \vdots \\ & & & 1 & \bar{\beta}_{i-1} \\ & & & & 1 \end{pmatrix}$

<u>*Remark:*</u> $\beta_i\,\beta_{i+1}\,\ldots\,\beta_j = \|r_j\|^2/\|r_{i-1}\|^2$, so entries of $\bar{\mathcal{B}}_i^{-1}$ may be arbitrarily large.

Analyzing pipelined CG by Ghysels et al.

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

$$\bar{x}_{i+1} = \bar{x}_i + \bar{\alpha}_i \bar{p}_i + \xi^x_{i+1}, \qquad \bar{s}_i = \bar{w}_i + \bar{\beta}_i \bar{s}_{i-1} + \xi^s_i,$$
$$\bar{r}_{i+1} = \bar{r}_i - \bar{\alpha}_i \bar{s}_i + \xi^r_{i+1}, \qquad \bar{w}_{i+1} = \bar{w}_i - \bar{\alpha}_i \bar{z}_i + \xi^w_{i+1},$$
$$\bar{p}_i = \bar{r}_i + \bar{\beta}_i \bar{p}_{i-1} + \xi^p_i, \qquad \bar{z}_i = A\bar{w}_i + \bar{\beta}_i \bar{z}_{i-1} + \xi^z_i,$$

The gap on the residual is coupled to the gaps on the auxiliary variables:

$$\begin{aligned}(B - A\bar{X}_i) - \bar{R}_i = {} & (A\Theta^{\bar{x}}_i + \Theta^{\bar{r}}_i)\, E_i + (A\Theta^{\bar{p}}_i + \Theta^{\bar{s}}_i)\, \bar{\mathcal{B}}_i^{-1} \bar{\mathcal{A}}_i \\ & + (A\Theta^{\bar{u}}_i + \Theta^{\bar{w}}_i)\, E_i \bar{\mathcal{B}}_i^{-1} \bar{\mathcal{A}}_i + (A\Theta^{\bar{q}}_i + \Theta^{\bar{z}}_i)\, \bar{\mathcal{B}}_i^{-1} \bar{\mathcal{A}}_i \bar{\mathcal{B}}_i^{-1} \bar{\mathcal{A}}_i\end{aligned}$$

with $$\bar{\mathcal{A}}_i = \begin{pmatrix} 0 & \bar{\alpha}_0 & \bar{\alpha}_0 & \cdots & \bar{\alpha}_0 \\ & 0 & \bar{\alpha}_1 & \cdots & \bar{\alpha}_1 \\ & & \ddots & & \vdots \\ & & & 0 & \bar{\alpha}_{i-2} \\ & & & & 0 \end{pmatrix}, \quad \bar{\mathcal{B}}_i^{-1} = \begin{pmatrix} 1 & \bar{\beta}_1 & \bar{\beta}_1\bar{\beta}_2 & \cdots & \bar{\beta}_1\bar{\beta}_2\ldots\bar{\beta}_{i-1} \\ & 1 & \bar{\beta}_2 & & \bar{\beta}_2\ldots\bar{\beta}_{i-1} \\ & & \ddots & & \vdots \\ & & & 1 & \bar{\beta}_{i-1} \\ & & & & 1 \end{pmatrix}$$

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

Carson et al. (2018) C. et al. (2018)

Analyzing deep ℓ-length pipelined CG

The recurrence relations for $\bar{x}_i$ and $\bar{p}_i$ in finite precision p(l)-CG are

$$\begin{aligned} \bar{x}_i &= \bar{x}_{i-1} + \bar{\zeta}_{i-1}\bar{p}_{i-1} + \xi_i^{\bar{x}} &\Leftrightarrow\quad \bar{x}_i &= \bar{x}_0 + \bar{P}_i\bar{q}_i + \Xi_i^{\bar{x}}\,\mathbf{1}, \\ \bar{p}_i &= (\bar{v}_i - \bar{\delta}_{i-1}\bar{p}_{i-1})/\bar{\eta}_i + \xi_i^{\bar{p}} &\Leftrightarrow\quad \bar{V}_i &= \bar{P}_i\bar{U}_i + \Xi_i^{\bar{p}}, \end{aligned}$$

with $\bar{T}_i = \bar{L}_i\bar{U}_i$, implying the actual residual equals

$$\begin{aligned} b - A\bar{x}_i &= \bar{r}_0 - A\bar{V}_i\bar{U}_i^{-1}\bar{q}_i + \overbrace{A\Xi_i^{\bar{p}}\bar{U}_i^{-1}\bar{q}_i - A\Xi_i^{\bar{x}}\,\mathbf{1} + \xi_0^{\bar{r}}}^{\text{Local Rounding Errors (LRE)}} \\ &= \bar{r}_0 - \bar{V}_{i+1}\bar{T}_{i+1,i}\bar{U}_i^{-1}\bar{q}_i - (A\bar{V}_i - \bar{V}_{i+1}\bar{T}_{i+1,i})\bar{U}_i^{-1}\bar{q}_i + \text{LRE} \\ &= \bar{r}_i - (A\bar{V}_i - \bar{V}_{i+1}\bar{T}_{i+1,i})\bar{U}_i^{-1}\bar{q}_i + \text{LRE} \end{aligned}$$

$\downarrow$ Computed residual **tends to zero**

$\downarrow$ **Inexact Lanczos relation** ("gap on $\bar{V}_{i+1}$") determines **maximal attainable accuracy**

Analyzing deep ℓ-length pipelined CG

Basis vector recurrences in finite precision p(l)-CG

$$\bar{v}_{i+1} = \left(\bar{z}_{i+1} - \sum_{j=i-2l+1}^{i} \bar{g}_{j,i+1}\bar{v}_i\right)/\bar{g}_{i+1,i+1} + \xi^{\bar{v}}_{i+1}, \quad \Leftrightarrow \quad \bar{Z}_i = \bar{V}_i\bar{G}_i + \Xi^{\bar{v}}_i \tag{1}$$

$$\bar{z}_{i+1} = (A\bar{z}_i - \bar{\gamma}_{i-l}\bar{z}_i - \bar{\delta}_{i-l-1}\bar{z}_{i-1})/\bar{\delta}_{i-l} + \xi^{\bar{z}}_{i+1}, \quad \Leftrightarrow A\bar{Z}_i = \bar{Z}_{i+1}\bar{B}_{i+1,i} + \Xi^{\bar{z}}_i \tag{2}$$

and the finite precision coefficient relation $\bar{G}_{i+1}\bar{B}_{i+1,i} = \bar{T}_{i+1,i}\bar{G}_i$ (3)
allow to compute the gap on the basis $\bar{V}_{i+1}$ as

$$\begin{aligned} A\bar{V}_i - \bar{V}_{i+1}\bar{T}_{i+1,i} &\overset{(1)}{=} A\bar{Z}_i\bar{G}_i^{-1} - \bar{Z}_{i+1}\bar{G}_{i+1}^{-1}\bar{T}_{i+1,i} - A\Xi^{\bar{v}}_i\bar{G}_i^{-1} + \Xi^{\bar{v}}_{i+1}\bar{G}_{i+1}^{-1}\bar{T}_{i+1,i} \\ &\overset{(3)}{=} (A\bar{Z}_i - \bar{Z}_{i+1}\bar{B}_{i+1,i} - A\Xi^{\bar{v}}_i + \Xi^{\bar{v}}_{i+1}\bar{B}_{i+1,i})\bar{G}_i^{-1} \\ &\overset{(2)}{=} (\Xi^{\bar{z}}_i - A\Xi^{\bar{v}}_i + \Xi^{\bar{v}}_{i+1}\bar{B}_{i+1,i})\,\bar{G}_i^{-1}. \end{aligned}$$

Amplification of local rounding errors possible, depending on $\bar{G}_i^{-1}$.

Cornelis et al. (2018)

Analyzing deep ℓ-length pipelined CG

Residual history $\|b - A\bar{x}_i\|_2$

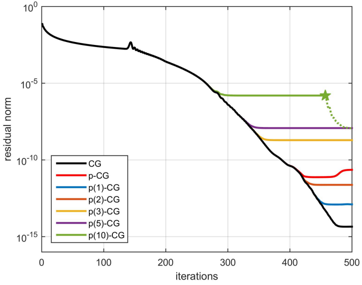

Norm $\|\bar{G}_i^{-1}\|_{\max}$

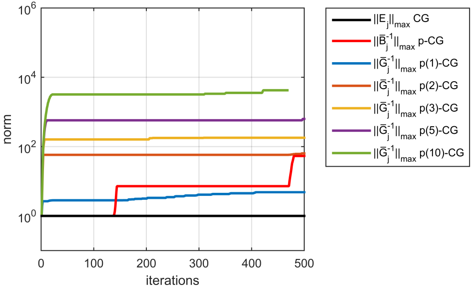

- The norm $\|\bar{G}_i^{-1}\|_{\max}$ quantifies the impact of rounding error amplification on attainable accuracy in p(l)-CG.
- The Cholesky factorization $Z_i^T Z_i = G_i^T G_i$ relates the conditioning of G_i and the auxiliary basis Z_i; numerical stability depends on the polynomial $P_l(A)$.

Hoemmen (2010) Ghysels et al. (2013)

Residual replacement in p-CG by Ghysels et al.

- Replace $\bar{r}_i = \mathrm{fl}(b - A\bar{x}_i), \bar{w}_i = \mathrm{fl}(A\bar{r}_i), \bar{s}_i = \mathrm{fl}(A\bar{p}_i), \bar{z}_i = \mathrm{fl}(A\bar{s}_i)$ in selected iterations
  Sleijpen et al. (1996) van der Vorst & Ye (2000) Strakos & Tichy (2002)
- Automated procedure based on estimate $\|b - A\bar{x}_i - \bar{r}_i\|$ (computed inexpensively)
  Carson & Demmel (2014) C. et al. (2018)
  - Replace sufficiently often such that residual gap remains small
  - Don't replace if $\|\bar{r}_i\|$ is small, which may cause delay of convergence

Speedup over single-node CG (12-240 cores)

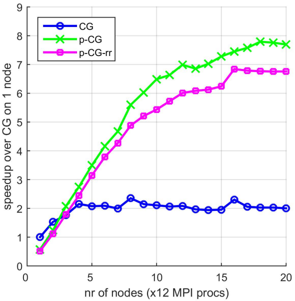

Accuracy vs. total time spent (240 cores)

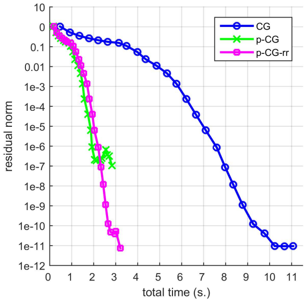

Stable recurrences for deep ℓ-length pipelined CG

Introduce l auxiliary bases

$$Z_{i+1}^{(0)} = [v_0, \dots v_i], \quad Z_{i+1}^{(1)} = [z_0^{(1)}, \dots z_i^{(1)}], \quad \dots \quad , \quad Z_{i+1}^{(l)} = [z_0, \dots z_i],$$

and replace the multi-term recurrence relation for v_{i-l+1} ($\sim 2l$ terms) by $l+1$ coupled three-term recurrence relations

$$\begin{cases} v_{i-l+1} &= (z_{i-l+1}^{(1)} + (\sigma_0 - \gamma_{i-l})v_{i-l} - \delta_{i-l-1}v_{i-l-1})/\delta_{i-l}, \\ z_{i-l+2}^{(1)} &= (z_{i-l+2}^{(2)} + (\sigma_1 - \gamma_{i-l})z_{i-l+1}^{(1)} - \delta_{i-l-1}z_{i-l}^{(1)})/\delta_{i-l}, \\ \quad\vdots & \qquad\qquad\qquad\qquad \vdots \\ z_i^{(l-1)} &= (z_i + (\sigma_{l-1} - \gamma_{i-l})z_{i-1}^{(l-1)} - \delta_{i-l-1}z_{i-2}^{(l-1)})/\delta_{i-l}, \\ z_{i+1} &= (Az_i - \gamma_{i-l}z_i - \delta_{i-l-1}z_{i-1})/\delta_{i-l}. \qquad \leftarrow 1 \text{ SpMV} \end{cases}$$

This modification causes (almost) no overhead

- the computational cost (#SpMVs and #axpy's) is identical to before,
- the storage cost increases by only $l-2$ vectors.

Stable recurrences for deep ℓ-length pipelined CG

IN FINITE PRECISION ARITHMETIC

Introduce l auxiliary bases

$$\bar{Z}_{i+1}^{(0)} = [\bar{v}_0, \dots \bar{v}_i], \quad \bar{Z}_{i+1}^{(1)} = [\bar{z}_0^{(1)}, \dots \bar{z}_i^{(1)}], \quad \dots \quad , \quad \bar{Z}_{i+1}^{(l)} = [\bar{z}_0, \dots \bar{z}_i],$$

and replace the multi-term recurrence relation for $\bar{v}_{i-l+1}$ ($\sim 2l$ terms) by $l+1$ coupled three-term recurrence relations that all introduce local rounding errors

$$\begin{cases} \bar{v}_{i-l+1} & = (\bar{z}_{i-l+1}^{(1)} + (\sigma_0 - \bar{\gamma}_{i-l})\bar{v}_{i-l} - \bar{\delta}_{i-l-1}\bar{v}_{i-l-1})/\bar{\delta}_{i-l} + \xi_{i-l+1}^{(0)}, \\ \bar{z}_{i-l+2}^{(1)} & = (\bar{z}_{i-l+2}^{(2)} + (\sigma_1 - \bar{\gamma}_{i-l})\bar{z}_{i-l+1}^{(1)} - \bar{\delta}_{i-l-1}\bar{z}_{i-l}^{(1)})/\bar{\delta}_{i-l} + \xi_{i-l+2}^{(1)}, \\ \quad \vdots & \qquad \qquad \vdots \\ \bar{z}_{i}^{(l-1)} & = (\bar{z}_i + (\sigma_{l-1} - \bar{\gamma}_{i-l})\bar{z}_{i-1}^{(l-1)} - \bar{\delta}_{i-l-1}\bar{z}_{i-2}^{(l-1)})/\bar{\delta}_{i-l} + \xi_{i}^{(l-1)}, \\ \bar{z}_{i+1} & = (A\bar{z}_i - \bar{\gamma}_{i-l}\bar{z}_i - \bar{\delta}_{i-l-1}\bar{z}_{i-1})/\bar{\delta}_{i-l} + \xi_{i+1}^{(l)}. \end{cases}$$

Stable recurrences for deep ℓ-length pipelined CG

IN FINITE PRECISION ARITHMETIC

Introduce l auxiliary bases

$$\bar{Z}^{(0)}_{i+1} = [\bar{v}_0, \dots \bar{v}_i], \quad \bar{Z}^{(1)}_{i+1} = [\bar{z}^{(1)}_0, \dots \bar{z}^{(1)}_i], \quad \dots \quad , \quad \bar{Z}^{(l)}_{i+1} = [\bar{z}_0, \dots \bar{z}_i],$$

and replace the multi-term recurrence relation for $\bar{v}_{i-l+1}$ ($\sim 2l$ terms) by $l+1$ coupled three-term recurrence relations that are written in matrix notation as

$$\begin{cases} \bar{Z}^{(1)}_{2:i-l+1} &= \bar{Z}^{(0)}_{i-l+1}\,\bar{T}_{i-l+1,i-l} - \sigma_0 \bar{Z}^{(0)}_{i-l} - \Xi^{(0)}_{i-l+1}\bar{\Delta}_{i-l+1,i-l}, \\ \bar{Z}^{(2)}_{2:i-l+2} &= \bar{Z}^{(1)}_{i-l+2}\,\bar{T}_{i-l+2,i-l+1} - \sigma_1 \bar{Z}^{(1)}_{i-l+1} - \Xi^{(1)}_{i-l+2}\bar{\Delta}_{i-l+2,i-l+1}, \\ \quad\vdots & \qquad\qquad\qquad\vdots \\ \bar{Z}^{(l)}_{2:i} &= \bar{Z}^{(l-1)}_{i}\,\bar{T}_{i,i-1} - \sigma_{l-1} \bar{Z}^{(l-1)}_{i-1} - \Xi^{(l-1)}_{i}\bar{\Delta}_{i,i-1}, \\ A\bar{Z}^{(l)}_{i} &= \bar{Z}^{(l)}_{i+1}\,\bar{T}_{i+1,i} - \Xi^{(l)}_{i+1}\bar{\Delta}_{i+1,i}. \end{cases}$$

Stable recurrences for deep ℓ-length pipelined CG

IN FINITE PRECISION ARITHMETIC

Introduce l auxiliary bases

$$\bar{Z}_{i+1}^{(0)} = [\bar{v}_0, \dots \bar{v}_i], \quad \bar{Z}_{i+1}^{(1)} = [\bar{z}_0^{(1)}, \dots \bar{z}_i^{(1)}], \quad \dots \quad , \quad \bar{Z}_{i+1}^{(l)} = [\bar{z}_0, \dots \bar{z}_i],$$

and replace the multi-term recurrence relation for $\bar{v}_{i-l+1}$ ($\sim 2l$ terms) by $l+1$ coupled three-term recurrence relations that are written in matrix notation as

$$\rightarrow \left\{ \begin{array}{ll} \bar{Z}_{2:i-l+1}^{(1)} & = \bar{Z}_{i-l+1}^{(0)} \bar{T}_{i-l+1,i-l} - \sigma_0 \bar{Z}_{i-l}^{(0)} - \Xi_{i-l+1}^{(0)} \bar{\Delta}_{i-l+1,i-l}, \\ \bar{Z}_{2:i-l+2}^{(2)} & = \bar{Z}_{i-l+2}^{(1)} \bar{T}_{i-l+2,i-l+1} - \sigma_1 \bar{Z}_{i-l+1}^{(1)} - \Xi_{i-l+2}^{(1)} \bar{\Delta}_{i-l+2,i-l+1}, \\ \vdots & \qquad \vdots \\ \bar{Z}_{2:i}^{(l)} & = \bar{Z}_{i}^{(l-1)} \bar{T}_{i,i-1} - \sigma_{l-1} \bar{Z}_{i-1}^{(l-1)} - \Xi_{i}^{(l-1)} \bar{\Delta}_{i,i-1}, \\ A\bar{Z}_{i}^{(l)} & = \bar{Z}_{i+1}^{(l)} \bar{T}_{i+1,i} - \Xi_{i+1}^{(l)} \bar{\Delta}_{i+1,i}. \end{array} \right.$$

For $\bar{Z}_{i+1}^{(l)}$ the gap is given by

$\bar{\Delta}_{i+1,i}$ diagonal matrix

$$A\bar{Z}_{i}^{(l)} - \bar{Z}_{i+1}^{(l)} \bar{T}_{i+1,i} = -\Xi_{i+1}^{(l)} \bar{\Delta}_{i+1,i}$$

Stable recurrences for deep ℓ-length pipelined CG

IN FINITE PRECISION ARITHMETIC

Introduce l auxiliary bases

$$\bar{Z}_{i+1}^{(0)} = [\bar{v}_0, \dots \bar{v}_i], \quad \bar{Z}_{i+1}^{(1)} = [\bar{z}_0^{(1)}, \dots \bar{z}_i^{(1)}], \quad \dots \quad , \quad \bar{Z}_{i+1}^{(l)} = [\bar{z}_0, \dots \bar{z}_i],$$

and replace the multi-term recurrence relation for $\bar{v}_{i-l+1}$ ($\sim 2l$ terms) by $l+1$ coupled three-term recurrence relations that are written in matrix notation as

$$\longrightarrow \left\{ \begin{array}{ll} \bar{Z}_{2:i-l+1}^{(1)} & = \bar{Z}_{i-l+1}^{(0)} \bar{T}_{i-l+1,i-l} - \sigma_0 \bar{Z}_{i-l}^{(0)} - \Xi_{i-l+1}^{(0)} \bar{\Delta}_{i-l+1,i-l}, \\ \bar{Z}_{2:i-l+2}^{(2)} & = \bar{Z}_{i-l+2}^{(1)} \bar{T}_{i-l+2,i-l+1} - \sigma_1 \bar{Z}_{i-l+1}^{(1)} - \Xi_{i-l+2}^{(1)} \bar{\Delta}_{i-l+2,i-l+1}, \\ \vdots & \qquad \vdots \\ \bar{Z}_{2:i}^{(l)} & = \bar{Z}_{i}^{(l-1)} \bar{T}_{i,i-1} - \sigma_{l-1} \bar{Z}_{i-1}^{(l-1)} - \Xi_{i}^{(l-1)} \bar{\Delta}_{i,i-1}, \\ A\bar{Z}_{i}^{(l)} & = \bar{Z}_{i+1}^{(l)} \bar{T}_{i+1,i} - \Xi_{i+1}^{(l)} \bar{\Delta}_{i+1,i}. \end{array} \right.$$

For $\bar{Z}_{i+1}^{(l-1)}$ the gap is given by

$\bar{\Delta}_{i+1,i}$ diagonal matrix

$$A\bar{Z}_{i}^{(l-1)} - \bar{Z}_{i+1}^{(l-1)} \bar{T}_{i+1,i} = (A\bar{Z}_{i}^{(l)} - \bar{Z}_{i+1}^{(l)} \bar{T}_{i+1,i}) \bar{\Delta}_{i,i}^{+} + \Xi_{i}^{(l)} - \Xi_{i+1}^{(l-1)} \bar{\Delta}_{i+1,i}$$

Stable recurrences for deep ℓ-length pipelined CG

IN FINITE PRECISION ARITHMETIC

Introduce l auxiliary bases

$$\bar{Z}_{i+1}^{(0)} = [\bar{v}_0, \dots \bar{v}_i], \quad \bar{Z}_{i+1}^{(1)} = [\bar{z}_0^{(1)}, \dots \bar{z}_i^{(1)}], \quad \dots \quad , \quad \bar{Z}_{i+1}^{(l)} = [\bar{z}_0, \dots \bar{z}_i],$$

and replace the multi-term recurrence relation for $\bar{v}_{i-l+1}$ ($\sim 2l$ terms) by $l+1$ coupled three-term recurrence relations that are written in matrix notation as

$$\rightarrow \begin{cases} \bar{Z}_{2:i-l+1}^{(1)} &= \bar{Z}_{i-l+1}^{(0)} \bar{T}_{i-l+1,i-l} - \sigma_0 \bar{Z}_{i-l}^{(0)} - \Xi_{i-l+1}^{(0)} \bar{\Delta}_{i-l+1,i-l}, \\ \bar{Z}_{2:i-l+2}^{(2)} &= \bar{Z}_{i-l+2}^{(1)} \bar{T}_{i-l+2,i-l+1} - \sigma_1 \bar{Z}_{i-l+1}^{(1)} - \Xi_{i-l+2}^{(1)} \bar{\Delta}_{i-l+2,i-l+1}, \\ \quad \vdots & \qquad \qquad \vdots \\ \bar{Z}_{2:i}^{(l)} &= \bar{Z}_{i}^{(l-1)} \bar{T}_{i,i-1} - \sigma_{l-1} \bar{Z}_{i-1}^{(l-1)} - \Xi_{i}^{(l-1)} \bar{\Delta}_{i,i-1}, \\ A\bar{Z}_{i}^{(l)} &= \bar{Z}_{i+1}^{(l)} \bar{T}_{i+1,i} - \Xi_{i+1}^{(l)} \bar{\Delta}_{i+1,i}. \end{cases}$$

For general $\bar{Z}_{i+1}^{(k)}$ the gap is given by $\qquad k \in \{0, 1, \dots, l-1\}$

$$A\bar{Z}_{i}^{(k)} - \bar{Z}_{i+1}^{(k)} \bar{T}_{i+1,i} = (A\bar{Z}_{i}^{(k+1)} - \bar{Z}_{i+1}^{(k+1)} \bar{T}_{i+1,i})\bar{\Delta}_{i,i}^{+} + \Xi_{i}^{(k+1)} - \Xi_{i+1}^{(k)} \bar{\Delta}_{i+1,i}$$

Stable recurrences for deep ℓ-length pipelined CG

IN FINITE PRECISION ARITHMETIC

Introduce l auxiliary bases

$$\bar{Z}_{i+1}^{(0)} = [\bar{v}_0, \ldots \bar{v}_i], \quad \bar{Z}_{i+1}^{(1)} = [\bar{z}_0^{(1)}, \ldots \bar{z}_i^{(1)}], \quad \ldots \quad , \quad \bar{Z}_{i+1}^{(l)} = [\bar{z}_0, \ldots \bar{z}_i],$$

and replace the multi-term recurrence relation for $\bar{v}_{i-l+1}$ ($\sim 2l$ terms) by $l+1$ coupled three-term recurrence relations that are written in matrix notation as

$$\longrightarrow \begin{cases} \bar{Z}_{2:i-l+1}^{(1)} &= \bar{Z}_{i-l+1}^{(0)} \bar{T}_{i-l+1,i-l} - \sigma_0 \bar{Z}_{i-l}^{(0)} - \Xi_{i-l+1}^{(0)} \bar{\Delta}_{i-l+1,i-l}, \\ \bar{Z}_{2:i-l+2}^{(2)} &= \bar{Z}_{i-l+2}^{(1)} \bar{T}_{i-l+2,i-l+1} - \sigma_1 \bar{Z}_{i-l+1}^{(1)} - \Xi_{i-l+2}^{(1)} \bar{\Delta}_{i-l+2,i-l+1}, \\ \quad\vdots & \qquad\qquad\qquad \vdots \\ \bar{Z}_{2:i}^{(l)} &= \bar{Z}_{i}^{(l-1)} \bar{T}_{i,i-1} - \sigma_{l-1} \bar{Z}_{i-1}^{(l-1)} - \Xi_{i}^{(l-1)} \bar{\Delta}_{i,i-1}, \\ A\bar{Z}_{i}^{(l)} &= \bar{Z}_{i+1}^{(l)} \bar{T}_{i+1,i} - \Xi_{i+1}^{(l)} \bar{\Delta}_{i+1,i}. \end{cases}$$

Accumulation of local rounding errors, but no amplification, similar to classic CG. The method thus attains the same accuracy as classic CG!

C. et al. (2019)

Stable recurrences for deep ℓ-length pipelined CG

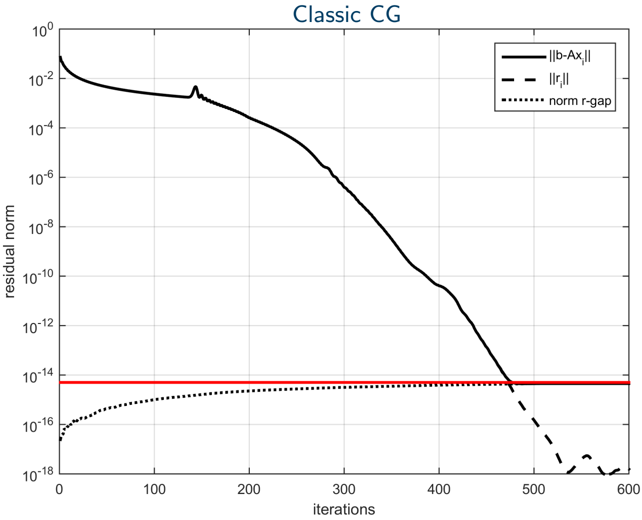

Stable recurrences for deep ℓ-length pipelined CG

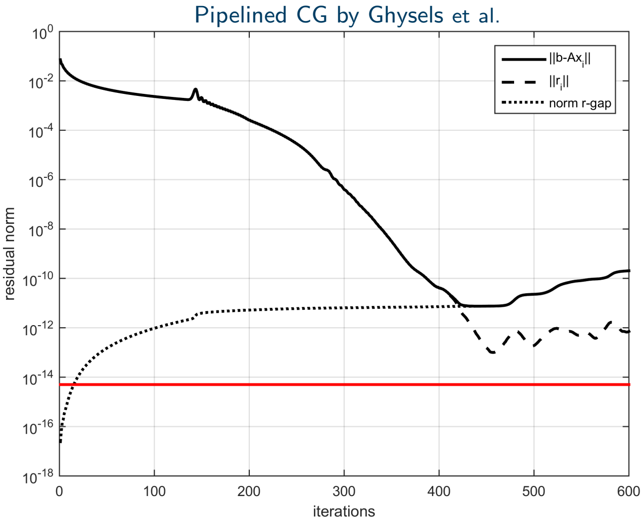

Stable recurrences for deep ℓ-length pipelined CG

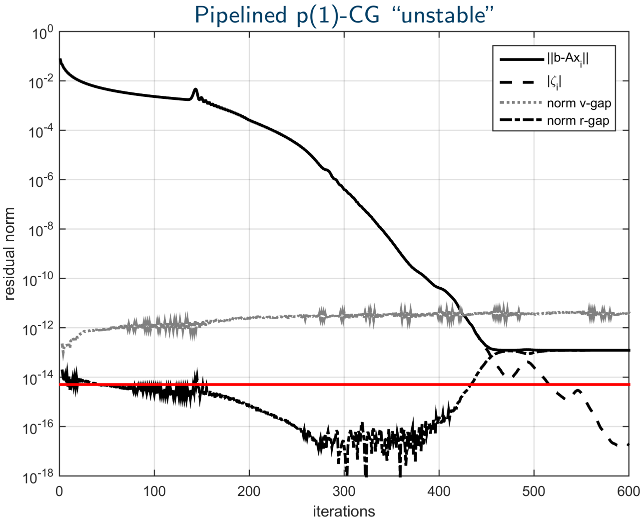

Stable recurrences for deep ℓ-length pipelined CG

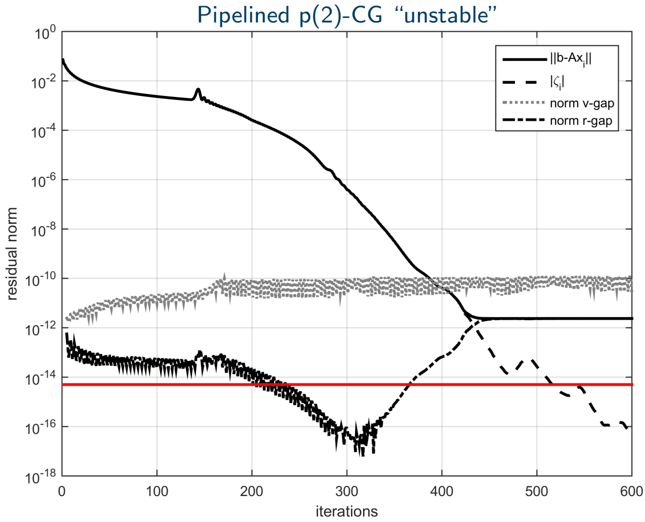

Stable recurrences for deep ℓ-length pipelined CG

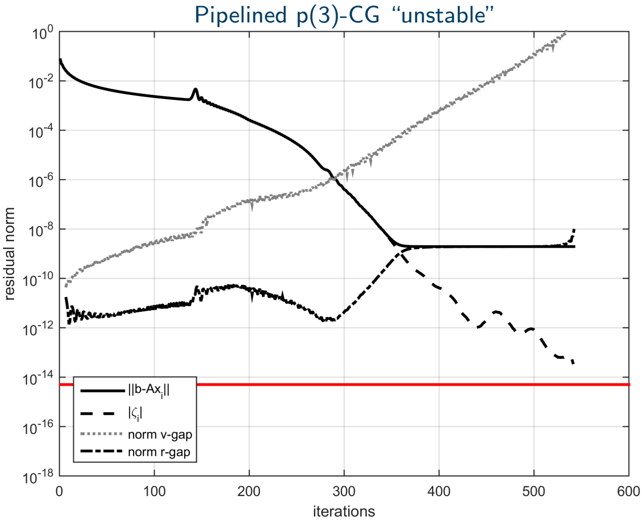

Stable recurrences for deep ℓ-length pipelined CG

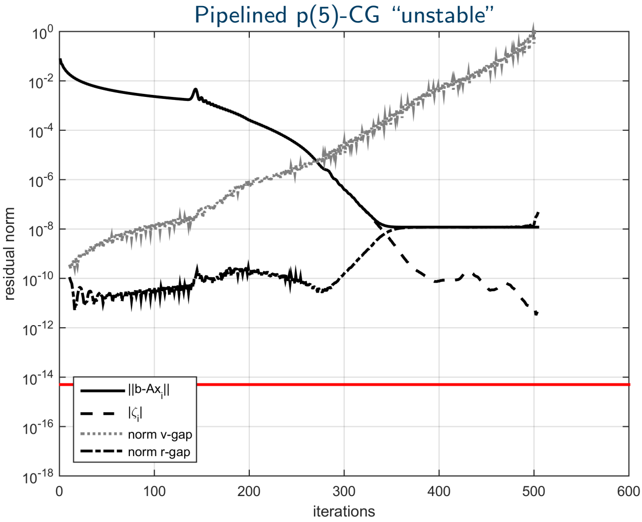

Stable recurrences for deep ℓ-length pipelined CG

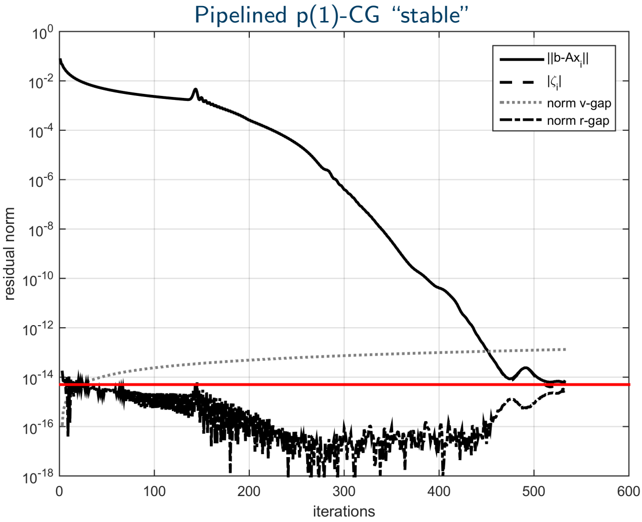

Stable recurrences for deep ℓ-length pipelined CG

Pipelined p(2)-CG "stable"

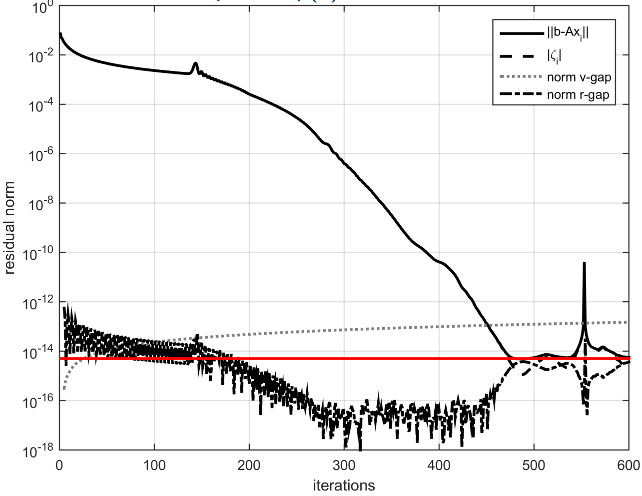

Stable recurrences for deep ℓ-length pipelined CG

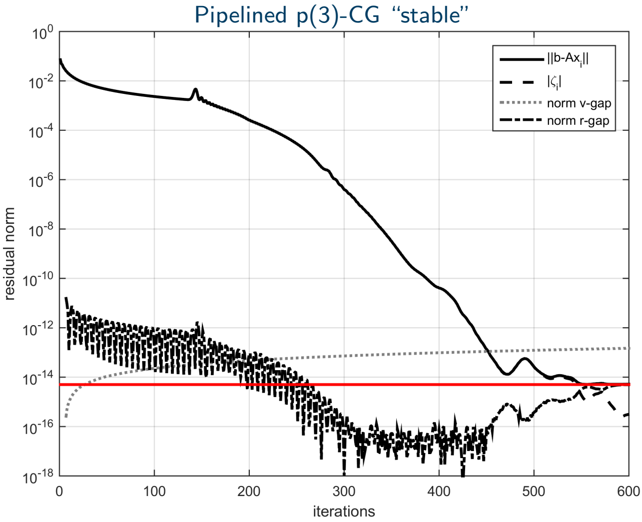

Stable recurrences for deep ℓ-length pipelined CG

Pipelined p(5)-CG "stable"

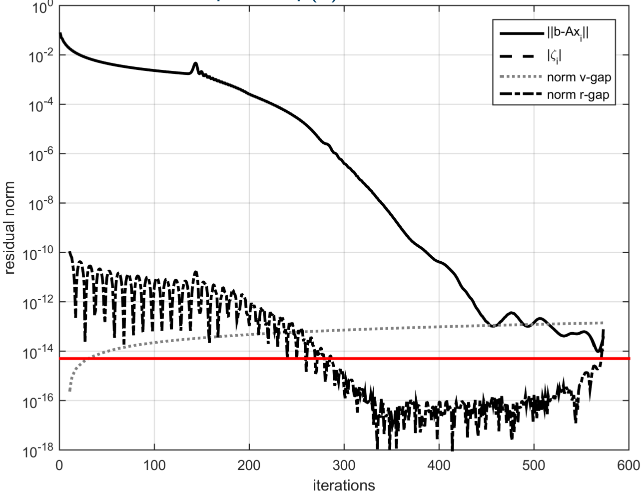

Deep ℓ-length pipelined CG

- Strong scaling on up to 32 14-core Intel E5-2680v4 Broadwell CPU nodes
- EDR Infiniband, Intel MPI 2018v3, PETSc v3.8.3, KSP ex2
- 2D 5-pt Poisson, 3 million unknowns, 1,500 iterations, no preconditioner

Speedup (over CG on 1 node)

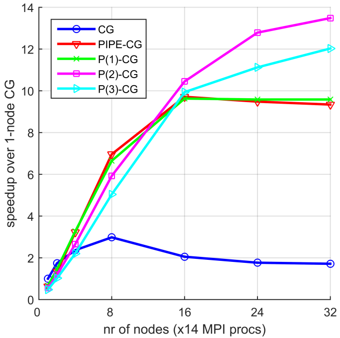

Accuracy (vs. total CPU time)

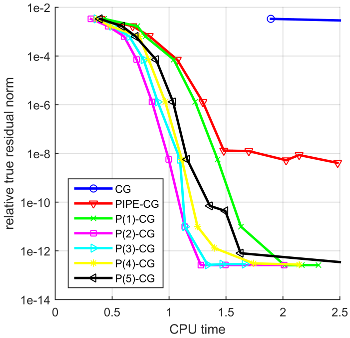

- Strong scaling on up to 128 14-core Intel E5-2680v4 Broadwell CPU nodes
- EDR Infiniband, Intel MPI 2018v3, PETSc v3.8.3, SNES ex48
- 3D Hydrostatic Ice Sheet Flow, 2.25 million FE, Newton-Krylov solver, 7 Newton steps, 4,500 total inner iter, block Jacobi preconditioner, inner tolerance: 1.0e-10, outer tolerance: 1.0e-8

Speedup (over CG on 1 node)

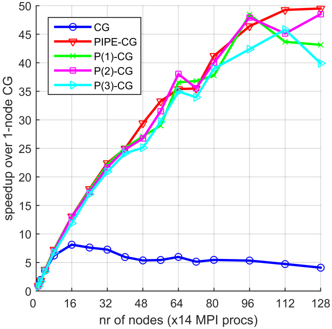

- Strong scaling on up to 128 14-core Intel E5-2680v4 Broadwell CPU nodes
- EDR Infiniband, Intel MPI 2018v3, PETSc v3.8.3, SNES ex48
- 3D Hydrostatic Ice Sheet Flow, 2.25 million FE, Newton-Krylov solver, 7 Newton steps, 4,500 total inner iter, block Jacobi preconditioner, inner tolerance: 1.0e-10, outer tolerance: 1.0e-8

Accuracy (vs. total number of inner iterations)

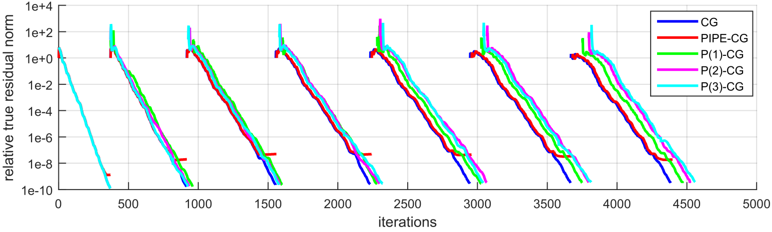

Takeaway messages

- Pipelined Krylov subspace methods are a promising approach
  - *Hide communication latency* behind computational kernels by adding auxiliary variables and recurrence relations
  - *p(ℓ)-CG: Deep pipelines* allow to hide global reduction phases behind multiple SpMV's/iterations
  - *Asynchronous implementation*: dot-products can take multiple iterations to complete; global reductions are implemented 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
- The finite precision behavior of communication avoiding- and hiding Krylov subspace algorithms should be carefully monitored
  - *Local rounding error analysis* allows to explain loss of attainable accuracy
- Insights to construct a more stable method are obtained from the analysis
  - *Fully restore attainable accuracy* in p(l)-CG at *no increase in computational costs or storage costs* through residual replacement-type techniques
  - The issue of *loss of orthogonality* has not been addressed by the modifications to p(l)-CG proposed in this talk

Contributions to PETSc

Open source HPC linear algebra toolkit: `https://www.mc`[illegible]`/petsc/`

We are soliciting for feedback from your applications

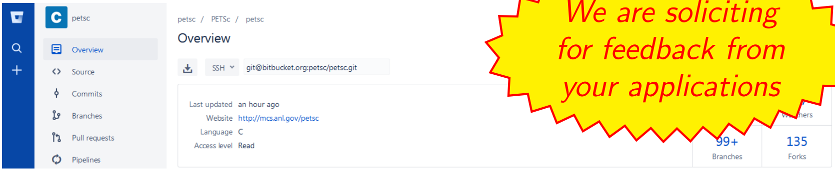

- KSPPGMRES: pipelined GMRES (thanks to J. Brown)
- KSPPIPECG: pipelined Conjugate Gradients
- KPPPIPECR: pipelined Conjugate Residuals
- KSPPIPECGRR: pipelined CG with automated residual replacement
- KSPPIPELCG: pipelined CG with deep pipelines
- KSPGROPPCG: asynchronous CG variant by W. Gropp and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

Related publications

Thank you!

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

jeffrey.cornelis@uantwerp.be
https://www.uantwerpen.be/en/staff/jeffrey-cornelis

wim.vanroose@uantwerp.be
https://www.uantwerpen.be/en/staff/wim-vanroose

- 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. C48–C71, 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.
- J. Cornelis, S. Cools, W. Vanroose, *The communication-hiding Conjugate Gradient method with deep pipelines.* Submitted to SIAM J. Sci. Comput., 2018, Preprint: ArXiv 1801.04728.
- S. Cools, J. Cornelis, W. Vanroose, *Numerically Stable Recurrence Relations for the Communication Hiding Pipelined Conjugate Gradient Method.* Submitted to IEEE Transactions on Parallel and Distributed Systems., 2019, Preprint: ArXiv 1902.03100.