Left-Preconditioned Communication-Avoiding Conjugate Gradient Methods for Multiphase CFD Simulations on the K Computer

> <u>Akie Mayumi</u>¹, Yasuhiro Idomura¹, Takuya Ina¹, Susumu Yamada¹, Toshiyuki Imamura² ¹Japan Atomic Energy Agency ²RIKEN

Acknowledgements

- T. Kawamura, S. Yamashita (JAEA)
- This work is supported by the MEXT, Grant for Post-K priority issue No. 6: Development of Innovative Clean Energy, and the computation is performed on the K-computer at the Riken (hp160208).





Exa-scale simulations for severe accident analysis

- JAEA promotes the development of multiphase thermal-hydraulic CFD code for analyzing severe accidents in the Fukushima Daiichi Nuclear Power Plant
- JUPITER code [Yamashita,ICONE2013] simulates relocation of molten materials in nuclear reactors as incompressible viscous fluids.

Peta-Scale

present JUPITER

- Finite difference method in structured grids
- Volume of fluid method for multiphase flows
- Multi-components (UO₂, Zry, B₄C, SUS)
- 3D domain decomposition (MPI+OpenMP)
- Target problems
 - Peta-scale (K-computer)
 - Simulate melt-relocation behavior of several fuel assemblies
 - Exa-scale
 - Severe accident analysis for whole reactor pressure vessel

Exa-Scale

Scalability issue of Poisson solver on K-computer

- K-computer(11.3PF/82,944nodes)
 - SPARC64VIIIfx(128GF=16GF x 8cores, B/F=0.5)
 - Tofu Interconnect(3D torus, 5GB/s x 4ways)
- Poisson solver in JUPITER



- 2nd order centered finite difference in structured grids (7-stencils)
- CG method with Block-Jacobi preconditioner (PETSc/in-house solver)
- Occupy over 90% of computational cost for large problems
- Large density contrast of multiphase flows gives ill-conditioned problem, and conditions become worse in large problems

→proper preconditioning is essential

- Scalability issue
 - P2P comm. scales on 3D torus
 - Collective comm. is bottleneck

→In this work, we resolve this issue using Communication Avoiding CG methods



Strong scaling of CG method on K-computer

Review of CA-CG method based on [Hoemmen, PhD2010]

CG

dependency between r_i and p_i

Algorithm 1 Conjugate Gradient methodInput: Ax = b, Initial guess x_1 Output: Approximate solution x_j 1: $r_1 := b - Ax_1, p_1 := r_1$ 2: for j = 1, 2, ... until convergence do3: Compute $w := Ap_j$ 4: $\alpha_j := \langle r_j, r_j \rangle / \langle w, p_j \rangle$ 5: $x_{j+1} := x_j + \alpha_j p_j$ 6: $r_{j+1} := r_j - \alpha_j w$ 7: $\beta_j := \langle r_{j+1}, r_{j+1} \rangle / \langle r_j, r_j \rangle$ 8: $p_{j+1} := r_{j+1} + \beta_j p_j$; test convergence

9: end for

Review of CA-CG method based on [Hoemmen, PhD2010]

CG

dependency between r_i and p_i

• CG with 3-term recurrence

use only r_i (decouple r_i-p_i dependency)

Algorithm 2 CG3 (3-term recurrence variant of CG) **Input:** Ax = b, Initial guess x_1 **Output:** Approximate solution x_i 1: $x_0 := 0, r_0 := 0, r_1 := b - Ax_1$ 2: for j = 1, 2, ... until convergence do 3: Compute $w_i := Ar_i$ P2P Comm. 4: $\mu_j := \langle r_j, r_j \rangle$; test for convergence 5: $\nu_j := \langle w_j, r_j \rangle$ Collective Comm. 6: $\gamma_i := \overline{\mu_i / \nu_i}$ 7: **if** j = 1 then 8: $\rho_i := 1$ else 9: $\rho_j := (1 - \frac{\gamma_j}{\gamma_{j-1}} \cdot \frac{\mu_j}{\mu_{j-1}} \cdot \frac{1}{\rho_{j-1}})^{-1}$ end if 10: 11: $x_{i+1} := \rho_i (x_i + \gamma_i r_i) + (1 - \rho_i) x_{i-1}$ 12: $r_{i+1} := \rho_i (r_i + \gamma_i w_i) + (1 - \rho_i) r_{i-1}$ 13: 14: end for

Review of CA-CG method based on [Hoemmen, PhD2010]

CG

dependency between r_i and p_i

CG with 3-term recurrence

use only r_i (decouple r_i-p_i dependency)

CG with outer/inner loops

```
Algorithm 3 CG3 with outer and inner loops
Input: Ax = b, Initial guess x_1
Output: Approximate solution x_i
 1: x_0 := 0, r_0 := 0, r_1 := b - Ax_1
 2: for k = 0, 1, 2, ... until convergence do
        for j = 1, 2, ..., s do
 3:
                                                     P2P Comm.
            Compute w_{sk+j} := Ar_{sk+j}
 4:
           \mu_{sk+j} := \langle r_{sk+j}, r_{sk+j} \rangle; test for convergence
 5:
           \nu_{sk+j} := \langle w_{sk+j}, r_{sk+j} \rangle
                                                     Collective Comm.
 6:
           \gamma_{sk+j} := \mu_{sk+j} / \nu_{sk+j}
 7:
           if sk + j = 1 then
 8:
              \rho_{sk+j} := 1
 9:
            else
10:
           \rho_{sk+j} := \left(1 - \frac{\gamma_{sk+j}}{\gamma_{sk+j-1}} \cdot \frac{\mu_{sk+j}}{\mu_{sk+j-1}} \cdot \frac{1}{\rho_{sk+j-1}}\right)^{-1}end if
11:
12:
           x_{sk+j+1} := \rho_{sk+j}(x_{sk+j} + \gamma_{sk+j}r_{sk+j}) + (1 - 1)
13:
            (\rho_{sk+j})x_{sk+j-1}
           r_{sk+j+1} := \rho_{sk+j}(r_{sk+j} + \gamma_{sk+j}w_{sk+j}) + (1 - 1)
14:
           (\rho_{sk+j})r_{sk+j-1}
        end for
15:
16: end for
```

Review of CA-CG method based on [Hoemmen, PhD2010]

CG

dependency between r_i and p_i

• CG with 3-term recurrence

use only r_i (decouple r_i-p_i dependency)

- CG with outer/inner loops
- CG with Matrix Power Kernel

```
compute Ar<sub>sk+j</sub> based on recurrence formula
using basis vectors given by MPK
```

 \rightarrow CA for halo comm.

Algorithm 4 CA-CG with MPK **Input:** Ax = b, Initial guess x_1 **Output:** Approximate solution x_{sk+i} 1: $x_0 := 0, r_0 := 0, r_1 := b - Ax_1$ 2: for k = 0, 1, 2, ... until convergence do $v_{sk+1} := r_{sk+1}$ 3: Compute MPK $\underline{V}_k = [v_{sk+1}, ..., v_{sk+s+1}]$ 4: for j = 1 to s do P2P Comm. 5: Compute d_{sk+j} $(Ar_{sk+j} = [R_{k-1}, \underline{V}_k]d_{sk+j})$ 6: $w_{sk+j} := [R_{k-1}, \underline{V}_k] d_{sk+j}$ 7: $\mu_{sk+j} := \langle r_{sk+j}, r_{sk+j} \rangle$; test for convergence 8: $\nu_{sk+j} := \langle w_{sk+j}, r_{sk+j} \rangle$ Collective Comm. 9: 10: $\gamma_{sk+j} := \mu_{sk+j} / \nu_{sk+j}$ if sk + j = 1 then 11: $\rho_{sk+j} := 1$ 12: 13: else $\rho_{sk+j} := \left(1 - \frac{\gamma_{sk+j}}{\gamma_{sk+j-1}} \cdot \frac{\mu_{sk+j}}{\mu_{sk+j-1}} \cdot \frac{1}{\rho_{sk+j-1}}\right)^{-1}$ 14: end if 15: $x_{sk+j+1} := \rho_{sk+j}(x_{sk+j} + \gamma_{sk+j}r_{sk+j}) + (1 - 1)$ 16: $(\rho_{sk+j})x_{sk+j-1}$ $r_{sk+j+1} := \rho_{sk+j}(r_{sk+j} + \gamma_{sk+j}w_{sk+j}) + (1 - 1)$ 17: $(\rho_{sk+j})r_{sk+j-1}$ end for 18: 19: end for

Review of CA-CG method based on [Hoemmen, PhD2010]

CG

dependency between r_i and p_i

CG with 3-term recurrence

use only r_i (decouple r_i-p_i dependency)

- CG with outer/inner loops
- CG with Matrix Power Kernel

compute Ar_{sk+j} based on recurrence formula using basis vectors given by MPK \rightarrow CA for halo comm.

CA-CG

compute inner product using Gram matrix→CA for reduction comm.

CA-CG is equivalent to CG in exact arithmetic

| Algorithm 5 CA-CO |
|--|
| Input: $Ax = b$, Initial guess x_1 |
| Output: Approximate solution x_{sk+j} |
| 1: $x_0 := 0, r_0 := 0, r_1 := b - Ax_1$ |
| 2: for $k = 0, 1, 2,$ until convergence do |
| 3: $v_{sk+1} := r_{sk+1}$ P2P Comm. |
| 4: Compute MPK $V_{k} = [v_{sk+1},, v_{sk+s+1}]$ |
| 5: $G_{k,k-1} := R_{k-1}^* \underline{V}_k, G_{kk} := \underline{V}_k^* \underline{V}_k$ Collective Comm. |
| 6: $G_k = \begin{pmatrix} D_{k-1} & G_{k,k-1} \\ G_{k,k-1}^* & G_{kk} \end{pmatrix}$ Gram Matrix |
| 7: for $j = 1$ to s do |
| 8: Compute d_{sk+j} $(Ar_{sk+j} = [R_{k-1}, \underline{V}_k]d_{sk+j})$ |
| 9: Compute g_{sk+j} $(r_{sk+j} = [R_{k-1}, \underline{V}_k]g_{sk+j})$ |
| 10: $w_{sk+j} := [R_{k-1}, \underline{V}_k] d_{sk+j}$ |
| 11: $\mu_{sk+j} := g^*_{sk+j} G_k g_{sk+j}$ |
| 12: $\nu_{sk+j} := g^*_{sk+j} G_k d_{sk+j}$ |
| 13: $\gamma_{sk+j} := \mu_{sk+j} / \nu_{sk+j}$ |
| 14: if $sk + j = 1$ then |
| 15: $\rho_{sk+j} := 1$ |
| 16: else |
| 17: $\rho_{sk+j} := \left(1 - \frac{\gamma_{sk+j}}{\gamma_{sk+j-1}} \cdot \frac{\mu_{sk+j}}{\mu_{sk+j-1}} \cdot \frac{1}{\rho_{sk+j-1}}\right)^{-1}$ |
| 18: end if |
| 19: $x_{sk+j+1} := \rho_{sk+j}(x_{sk+j} + \gamma_{sk+j}r_{sk+j}) + (1 - 1)$ |
| $ \rho_{sk+j})x_{sk+j-1} $ |
| 20: $r_{sk+j+1} := \rho_{sk+j}(r_{sk+j} + \gamma_{sk+j}w_{sk+j}) + (1 - 1)$ |
| $ ho_{sk+j})r_{sk+j-1}$ |
| 21: end for |
| 22: end for |

Algorithm 5 CA CG

Related works – Application of CA-Krylov methods

- Stability and convergence properties of CA-Krylov methods [Carson, PhD2015]
 - Convergence properties with different basis vectors
 - Improved convergence with residual replacement technique
- CA preconditioners for CA-Krylov methods [Yamazaki,SC14]
 - CA-GMRES implementation on GPUs
 - CA-preconditioning underlap approach
- Chebyshev basis CA-CG on K-computer [Kumagai, PPAM2015]
 - Strong scaling of CBCG up to 100k cores
- →Most of former works were successful for
 - Iarge CA steps with s>10
 - No or approximate preconditioning

→We apply LP-CA-CG to ill-conditioned problem which is limited to s=3



Convergence issue of CA-CG in JUPITER

Poor convergence of CA-CG in JUPITER

- CA steps with s>3 do not converge
- Causes of convergence degradation
 - Orthogonality of basis vectors generated during CA steps
 - Round off errors in inner product operations using Gram matrix
- Possible solutions
 - Newton and Chebyshev basis vectors [Hoemmen,PhD2010, Carson,PhD2015]
 - Mixed precision approach (quadruple precision only in Gram matrix)
 →Preliminary tests did not show large performance gain

 \rightarrow We pursue performance improvement at s=3



Convergence of CA-CG (JUPITER:800x500x3540)



Convergence of CA-CG with mixed precision (JUPITER:104x104x265)

Optimization of CA-CG on K-computer

- Serial optimization
 - Avoid indirect data access by changing data format from CSR to CDR
 - Minimize memory access by data blocking and loop splitting
 Maximize impact of CA by minimizing cost of calculation
- Comparisons of two Block-Jacobi based preconditioners
 - Original Block-Jacobi (bj) preconditioning
 - CA-preconditioning with underlap (u) approach [Yamazaki,SC14]
 - \rightarrow Explore preconditioners suitable for CA-CG on K-computer

Optimize data format for JUPITER

- Compressed Sparse Row (CSR) format
 - Widely used in many matrix libraries such as PETSc
 - Indirect memory access is overhead for structured grid data
- Compressed Diagonal Storage (CDS) format
 - Used in our in-house solvers
 - Direct memory access for structured grid data
 - \rightarrow 3.3x performance gain compared to CSR format



CDS format Offset from diagonal

Coefficients

SpMV sample code

 $\label{eq:for(j=0;j<ndia;j++)} offset=index[j]; \\for(i=0;i<n;i++) \{ \\ q[i]=q[i]+a[i+j*n]*p[i+offset]; \} \}$

Analysis of arithmetic intensity of CA-CG kernels

Outer loop: SpMV and Block-Jacobi preconditioning (62%)

| | SpMV | BJ-precond. |
|----------------------------|-------|-------------|
| Arithmetic intensity (f/b) | 0.163 | 0.116 |
| Roofline (Gflops) | 7.15 | 5.22 |
| Sustained (Gflops) | 6.37 | 4.99 |

 2^{nd} order centered finite difference with 7-stencils \rightarrow low arithmetic intensity

Outer loop: Gram matrix computation (10%)

| | s=1 | s=2 | s=3 |
|----------------------------|-------|-------|-------|
| Arithmetic intensity (f/b) | 0.300 | 0.469 | 0.636 |
| Roofline (Gflops) | 12.60 | 18.66 | 24.08 |
| Sustained (Gflops) | 14.35 | 23.60 | 27.16 |

$$\frac{f}{b} = \frac{(s+1) \times (2s+1) \times 2}{(3s+3) \times 8} \propto \mathcal{O}(s)$$

Inner loop: Inner product and 3-term recurrence (28%)

| | s=1 | s=2 | s=3 | | s=1 | s=2 | s=3 |
|--|-------|-------|---|---------------------------------|----------------------------|----------------------------|-------|
| Arithmetic intensity (f/b) | 0.105 | 0.118 | 0.130 | optimized | 0.188 | 0.354 | 0.521 |
| Roofline (Gflops) | 4.39 | 5.13 | 5.70 | | 8.18 | 14.62 | 20.40 |
| Sustained (Gflops) | 3.33 | 3.67 | 4.32 | | 7.30 | 12.85 | 19.14 |
| $\frac{f}{b} = \frac{((4s+1)\times 2+15)\times s+2}{((22+(2s+3)\times 2)\times 8)\times s} \propto const.$ | | | $\frac{f}{b} = \frac{((4s+1))}{((2s+1))}$ | $(-1) \times 2 + (+6) \times 2$ | $15) \times s +$ +2s)×8 | $\frac{2}{3} \propto O(s)$ | |

Improve arithmetic intensity of inner loop

Original

for j=1, s, j++ Compute \mathbf{d}_{j} , μ_{j} , ν_{j} , ρ_{j} , γ_{j} for i=1, n, i++ for k=1, 2s+1, k++ $\mathbf{u}_{j}[i] = \{\mathbf{Q}, \mathbf{V}\}[i, k]\mathbf{d}_{j}[k]$ $\mathbf{y}_{j}[i] = \{\mathbf{Z}, \mathbf{W}\}[i, k]\mathbf{d}_{j}[k]$ endfor $\mathbf{x}_{j+1}[i] = \rho_{j}(\mathbf{x}_{j}[i] + \gamma_{j}\mathbf{q}_{j}[i]) + (1-\rho_{j}) \mathbf{x}_{j-1}[i]$ $\mathbf{q}_{j+1}[i] = \rho_{j}(\mathbf{q}_{j}[i] - \gamma_{j}\mathbf{u}_{j}[i]) + (1-\rho_{j}) \mathbf{q}_{j-1}[i]$ $\mathbf{z}_{j+1}[i] = \rho_{j}(\mathbf{z}_{j}[i] - \gamma_{j}\mathbf{y}_{j}[i]) + (1-\rho_{j}) \mathbf{z}_{j-1}[i]$ ond for

endfor

endfor

{Q,V} and {Z,W} are loaded s times

Reuse {Q,V} and {Z,W} for s times via loop splitting and data blocking

 \rightarrow Reduce memory access from s² to s

Optimized

```
for j=1, s, j++
    Compute \mathbf{d}_i, \mu_i, \nu_j, \rho_j, \gamma_j
endfor
for ii=1 ,n ,ii=ii+nblock
    for j=1, s, j++
         for i=ii ,ii+nblock-1 ,i++
             for k=1, 2s+1, k++
                  \mathbf{u}_{i}[i] = \{\mathbf{Q}, \mathbf{V}\}[i, k]\mathbf{d}_{i}[k]
                  \mathbf{y}_{i}[i] = \{\mathbf{Z}, \mathbf{W}\}[i, k]\mathbf{d}_{i}[k]
             endfor
         endfor
    endfor
    for j=1, s, j++
         for i=ii ,ii+nblock-1 ,i++
             \mathbf{x}_{i+1}[i] = \rho_i(\mathbf{x}_i[i] + \gamma_i \mathbf{q}_i[i]) + (1 - \rho_i) \mathbf{x}_{i-1}[i]
             \mathbf{q}_{i+1}[i] = \rho_i(\mathbf{q}_i[i] - \gamma_i \mathbf{u}_i[i]) + (1 - \rho_i) \mathbf{q}_{i-1}[i]
             \mathbf{z}_{i+1}[i] = \rho_i(\mathbf{z}_i[i] - \gamma_i \mathbf{y}_i[i]) + (1 - \rho_i) \mathbf{z}_{i-1}[i]
         endfor
    endfor
endfor
```

Block-Jacobi based preconditioners for CA-CG

- Original Block-Jacobi preconditioning with SpMV
 - P2P comm. of halo data at every step
 - No additional comm. for BJ preconditioning
 - →Hybrid CA-CG method (CA inner product + no-CA SpMV)
- Block-Jacobi preconditioning with CA-SpMV
 - Halo data for s-steps is transferred in advance
 - Additional computation for extended halo data
 Additional P2P comm. for BJ preconditioning
- Underlap preconditioning with CA-SpMV
 - Point-Jacobi preconditioning for surface part
 - \rightarrow CA step and parallelization affect convergence

→Explore optimum preconditioner for K-computer



Block-Jacobi with CA-SpMV(s=2)





Converge property

- Matrix data from Poisson solver in JUPITER
 - Melt relocation of a fuel assembly
 - Problem size: n = 800 x 500 x 3,540 = 1.4 x 10⁹
- Comparisons of CG solvers
 - Original CG
 - Hybrid CA-CG with Block-Jacobi (bj)
 - CA-CG with underlap (u)

| CA steps | s=1 | s=2 | s=3 | s=4 |
|-----------|------|------|------|------|
| CG(bj) | 6206 | - | - | - |
| CA-CG(bj) | 6208 | 6224 | 6222 | × |
| CA-CG(u) | 6214 | 7188 | 7668 | 8072 |

CA step scan of number of iteration at 125 nodes

Node number scan of number of iteration at s=3

| Nodes | 1000 | 2000 | 4000 | 8000 |
|-----------|------|------|------|-------|
| CG(bj) | 6333 | 6428 | 6313 | 6633 |
| CA-CG(bj) | 6354 | 6444 | 6516 | 6636 |
| CA-CG(u) | 8274 | 8556 | 9582 | 10206 |

- Impact of CA step and parallelization on Block-Jacobi is weak
- Convergence of underlap degrades with CA step and number of nodes because of expansion of surface part with Point-Jacobi approximation



Strong scaling of CA-CG on K-computer



- CA-CG solvers show good strong scaling up to 30,000 nodes (240k cores)
- Hybrid CA-CG(bj) is fastest, and total cost is reduced by 47% from CG(bj)
- In CA-CG(u), advantage of CA is almost cancelled by worse convergence

Detailed cost distribution in CG and hybrid CA-CG



- Collective communication is reduced by 1/s=1/3
- P2P communication is almost comparable
- Calculation part shows different features between 15k and 30k nodes
 - Up to 15k nodes, calculation cost of CA-CG is slightly higher than CG
 - FP operations/Performance of CA-CG is increased by 2.4x/2x
 - At 30k nodes, calculation cost of CG is lower than CA-CG
 - Inner most loop size is not enough for pipelining

→Compute rich CA-CG can keep performance in such strong scaling limit

Conclusion

- LP-CA-CG method is applied to ill-conditioned pressure Poisson equation in multiphase CFD code JUPITER
 - CA procedures affect convergence and CA steps are limited to s=3
 - Even at s=3, significant performance gain was achieved on K-computer
 - Optimum data format for structured grid data
 - Improved arithmetic intensity with loop splitting implementation
 - Hybrid CA-CG approach based on 3D torus network

→CA-Krylov solver design strongly depends on computing platforms

Future Work

- Development of CA-Krylov solvers for different computing platforms
- Improvement of CA-Krylov methods

(e.g. preconditioning, basis vectors, residual replacement, mixed-precision)