Efficient Parallel Multigrid Methods on Manycore Clusters with Double/Single Precision Computing

Kengo Nakajima[*1,*2], **Takeshi Ogita**[*3]**, Masatoshi Kawai**[*1]
***1: Information Technology Center, University of Tokyo, *2: RIKEN R-CCS,**
***3: Tokyo Women's Christian University**

Efficient Parallel Multigrid Methods on Manycore Clusters with Double/Single Precision Computing

Kengo Nakajima[*1,*2], **Takeshi Ogita**[*3], **Masatoshi Kawai**[*1]
*1: Information Technology Center, University of Tokyo, *2: RIKEN R-CCS,
*3: Tokyo Women's Christian University

Acknowledgements

- JST/CREST
- DFG/SPPEXA
- JSPS Grant-in-Aid for Scientific Research (S): 19H05662
- JHPCN (jh200036-MDHI, jh200037-NAH, jh200041-NAH)
- JCAHPC
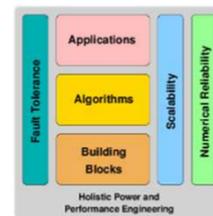
- Background & Overview
  - Multigtrid Methods
  - Overview of Previous/Present Works
- SELL-C-σ with Double Precision Computing
- Computing in Double/Single Precision
- Accuracy Verification
- Summary

Features of Multigrid (MG) Methods

✓ **Scalable multilevel method for solving linear equations**

✓ GMG (Geometrical Multigrid) and AMG (Algebraic Multigrid)

✓ **Number of iterations until convergence for multigrid method is kept constant as the problem size changes, Comp. Time = O(N)**

✓ The parallel multigrid method is expected to be one of the most powerful tools on exa-scale systems.

✓ Applied to rather well-conditioned problems (e.g. Poisson's eqn's)

  ❑ Many sophisticated methods for real-world applications are under development (next presentation)

✓ **MG is scalable, but there are many things to be done towards exascale computing**

smoothing (relaxation)

prolongation (interpolation)

restriction

Error on the fine grid

Error approximated on a smaller coarse grid

[LLNL]
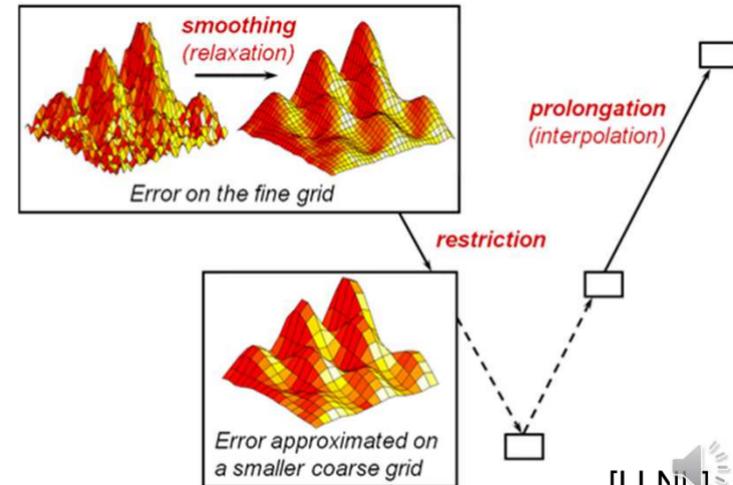
Overview (1/3)

- The parallel multigrid method is expected to play an important role in scientific computing on exa-scale supercomputer systems for solving large-scale linear equations with sparse coefficient matrices.

- Because solving sparse linear systems is a very memory-bound process, efficient method for storage of coefficient matrices is a crucial issue.

- In the previous works, authors implemented **Sliced ELL** method to parallel conjugate gradient solvers with multigrid preconditioning (**MGCG**) for the application on 3D groundwater flow through heterogeneous porous media (pGW3D-FVM), and excellent performance has been obtained on large-scale multicore/manycore clusters.

pGW3D-FVM: Target Application
[KN IEEE ICPADS 2014] (Best Paper Award)

- 3D Groundwater Flow via Heterogeneous Porous Media
- Finite-Volume Method on Structured Cubic Voxel Mesh
- Poisson's equation $\nabla \cdot \left(\lambda(x, y, z) \nabla \phi \right) = q$
  - Randomly distributed water conductivity (λ)
  - $\lambda = 10^{-5} \sim 10^{+5}$, Average: 1.00
- **Conjugate Gradient preconditioned by Multigrid (MGCG)**
  - **Geometric Multigrid (GMG): Octree-based**
  - **IC(0) Smoother, V-Cycle**
  - **Additive Schwartz Domain Decomposition**
- **Sliced ELL for Storage of Sparse Matrices**
- **Fortran90, MPI/OpenMP**
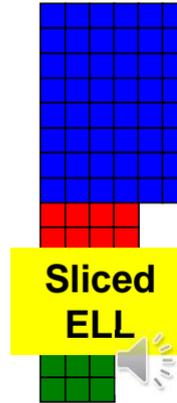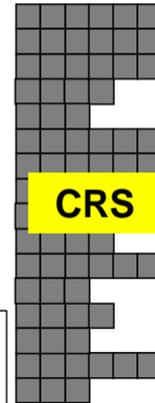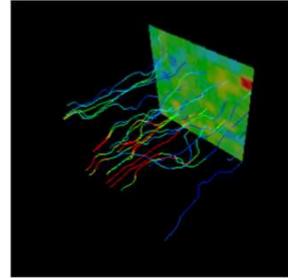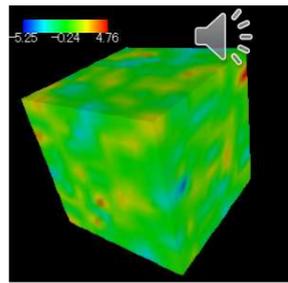
CRS **ELL** **Sliced ELL**

Nakajima, K., Optimization of Serial and Parallel Communications for Parallel Geometric Multigrid Method, Proceedings of the 20th IEEE International Conference for Parallel and Distributed Systems (ICPADS 2014) 25-32, Hsin-Chu, Taiwan, 2014

Effects of Sliced-ELL on MGCG/pGW3D-FVM
[KN ICPADS 2014]

- Fujitsu PRIMEHPC FX10
  - Weak-Scalingup to 4,096-nodes, 655,536-cores
  - max 17,179,869,184 DOF
  - HB 8x2
- 1.9x performance@655,536-cores: (Sliced ELL+CGA) over CRS
- MGCG: Sparse Linear Solver
  - Typical Memory-Bound Procedure
  - Effects of Memory Access/Matrix Storage are significant

Legend: ● CRS, △ ELL, ▲ ELL+CGA, ▲ Sliced ELL+CGA

y-axis: sec. (0.0 to 20.0); x-axis: CORE# (100 to 100000)

Very Scalable for Weak Scaling

Nakajima, K., Optimization of Serial and Parallel Communications for Parallel Geometric Multigrid Method, Proceedings of the 20th IEEE International Conference for Parallel and Distributed Systems (ICPADS 2014) 25-32, Hsin-Chu, Taiwan, 2014
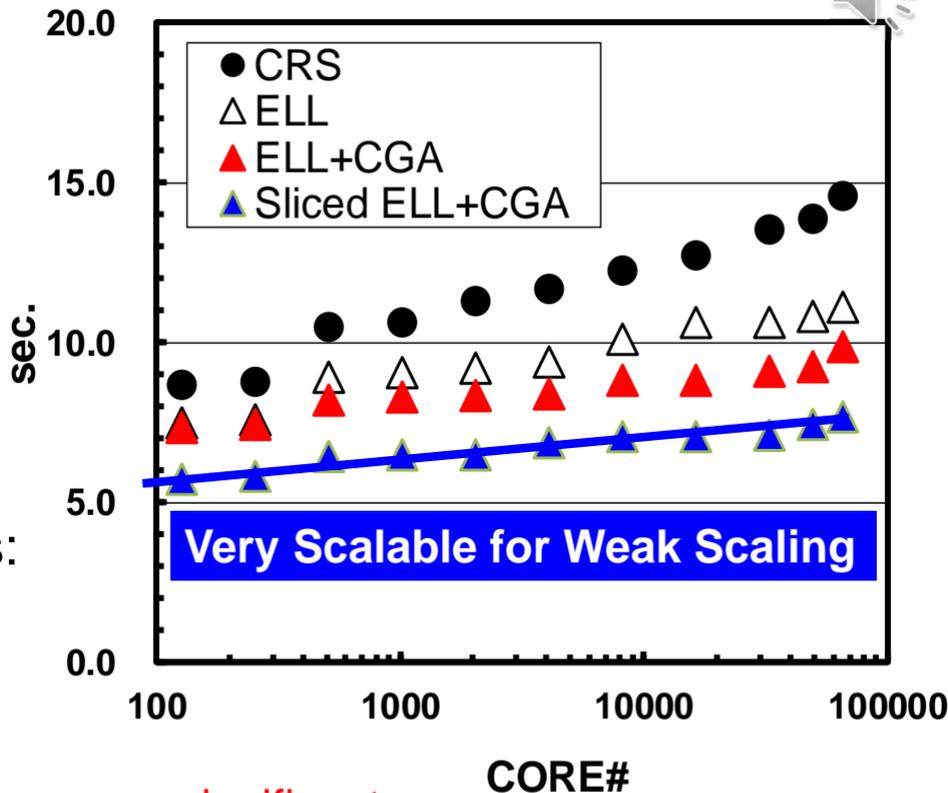
Overview (2/3)

- In the present work, authors introduced SELL-C-σ with double/single precision computing to the MGCG solver, and evaluated the performance of the solver with OpenMP/MPI hybrid parallel programing models on the Oakforest-PACS (OFP) system at JCAHPC **using up to 2,048 nodes of Intel Xeon Phi**.

- **Because SELL-C-σ is suitable for wide-SIMD architecture, such as Xeon Phi, improvement of the performance over the sliced ELL was more than 35% for double precision and more than 45% for single precision on OFP.**

- **Finally, accuracy verification was conducted based on the method proposed by authors for solving linear equations with sparse coefficient matrices with M-property.**

Overview (3/3)

- **This is one of the first examples of SELL-C-σ applied to forward/backward substitutions in ILU-type smoother of multigrid solver with double/single precision computing.**
- **The effect of SELL-C-σ for computing with single precision (FP32) is very significant.**

- Background & Overview
  - Multigtrid Methods
  - Overview of Previous/Present Works
- **SELL-C-σ with Double Precision Computing**
- Computing in Double/Single Precision
- Accuracy Verification
- Summary

Target System: Oakforest-PACS (OFP)

- **Intel Xeon Phi (Knights Landing, KNL), OPA**, Fujitsu
- 8,208 nodes, 25+PF, 22nd in TOP 500 (November 2020)
- Operated by JCAHPC (U.Tsukuba & U.Tokyo)
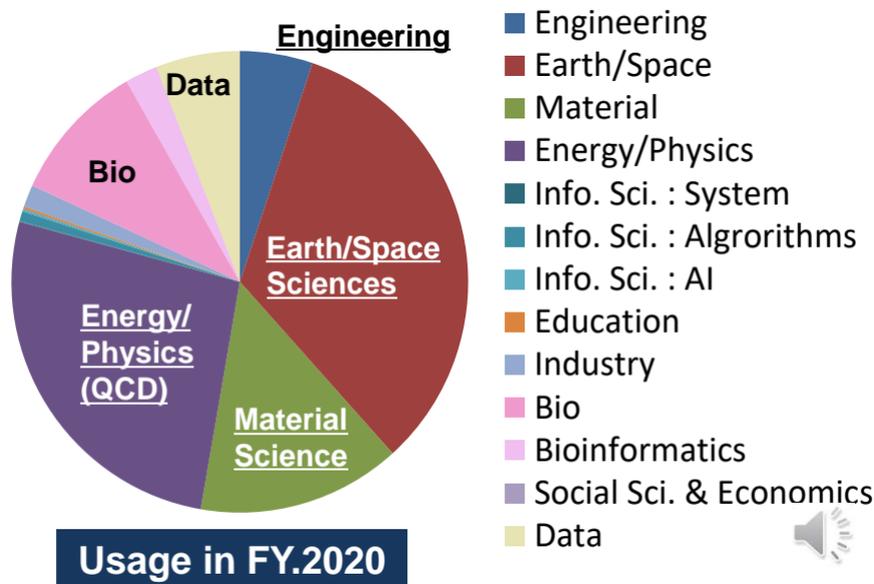
Usage in FY.2020

Pie chart legend:
- Engineering
- Earth/Space
- Material
- Energy/Physics
- Info. Sci. : System
- Info. Sci. : Algrorithms
- Info. Sci. : AI
- Education
- Industry
- Bio
- Bioinformatics
- Social Sci. & Economics
- Data

Pie chart labels: Engineering, Data, Bio, Earth/Space Sciences, Energy/Physics (QCD), Material Science

Main Purpose of the Present Work: SELL-C-σ

- **Applying SELL-C-σ to MGCG Solvers in pGW3D-FVM**
  - **Currently with Sliced ELL**

- Sliced ELL and SELL-C-σ have been mostly applied to SpMV, or Gauss-Seidel Iterative Solvers/Smoothers

- Implementation to Forward/Backward Substitution in ILU-type Smoothers is very difficult
  - First example of Sliced ELL [KN IEEE ICPADS 2014]
  - This is the first example of SELL-C-σ
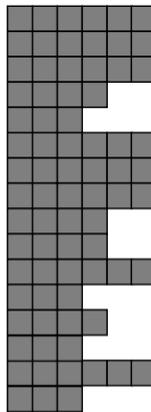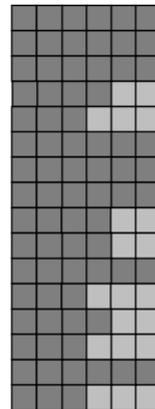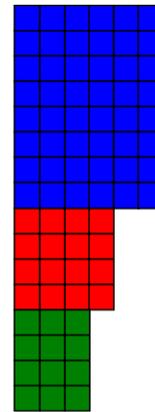
CRS   ELL   Sliced ELL   SELL-C-σ
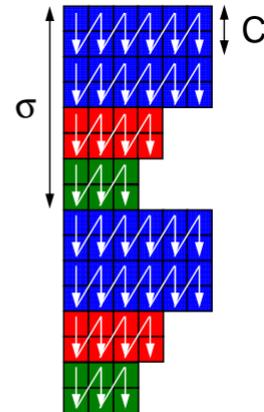
Constructing SELL-C-σ

1. Pick chunk size C (guided by SIMD/T widths)
2. Pick sorting scope σ
3. Sort rows by length within each sorting scope
4. Pad chunks with zeros to make them rectangular
5. Store matrix data in "chunk column major order"

"Chunk occupancy": fraction of "useful" matrix entries

$$\beta = \frac{N_{nz}}{\sum_{i=0}^{N_c} C \cdot l_i}$$

Width of chunk i: l_i

Sorting scope σ

Chunk size C

sorted

sorted
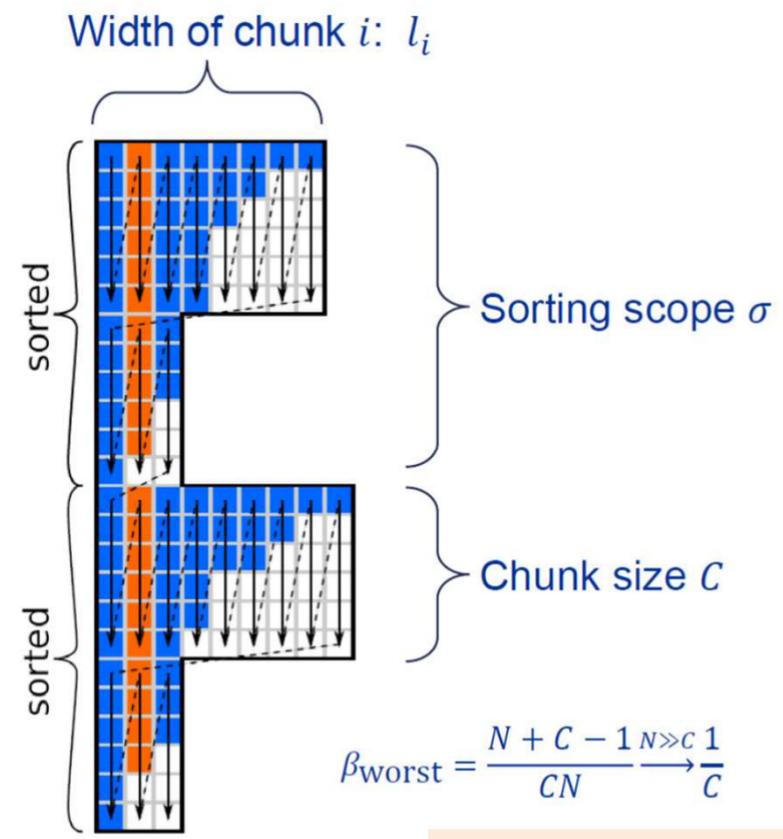
$$\beta_{worst} = \frac{N + C - 1}{CN} \xrightarrow{N \gg C} \frac{1}{C}$$

SELL-6-12
β=0.66

[Kreutzer, Hager, Wellein, SIAM SISC 2014]

Main Purpose of the Present Work: SELL-C-σ

- Applying SELL-C-σ to MGCG Solvers in pGW3D-FVM
  - Currently with Sliced ELL

- Sliced ELL and SELL-C-σ have been mostly applied to SpMV, or Gauss-Seidel Iterative Solvers/Smoothers

- **Implementation to Forward/Backward Substitution in ILU-type Smoothers is very difficult**
  - **First example of Sliced ELL [KN IEEE ICPADS 2014]**
  - **This is the first example of SELL-C-σ**
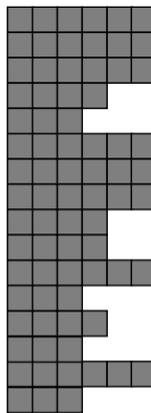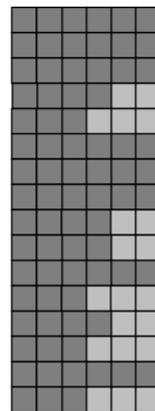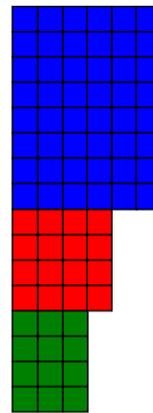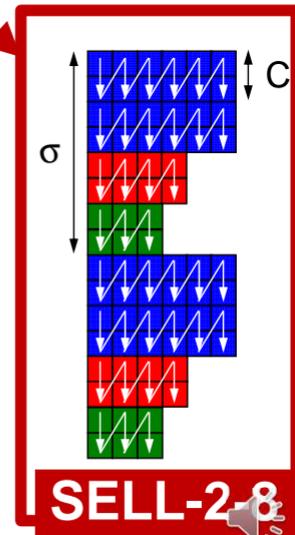
CRS   ELL   Sliced ELL   SELL-2-8

Coalesced, CM-RCM(2)

- CM-RCM with 2 colors: CM-RCM(2)
  - Number of iterations will increase with 2-colors, compared to RCM (current method)
  - Performance on OFP with multithreading is better with fewer colors
    - Loop length
  - Implementation of SELL-C-σ is easier than RCM
- Coloring part is not parallelized, but implementation of CM-RCM(2) is easy

Red-Black, 2色のMC

61	29	62	30	63	31	64	32
25	57	26	58	27	59	28	60
53	21	54	22	55	23	56	24
17	49	18	50	19	51	20	52
45	13	46	14	47	15	48	16
9	41	10	42	11	43	12	44
37	5	38	6	39	7	40	8
1	33	2	34	3	35	4	36

CM-RCM（2）, Coalesced

13	42	7	37	3	34	1	33
49	14	43	8	38	4	35	2
21	50	15	44	9	39	5	36
56	22	51	16	45	10	40	6
27	57	23	52	17	46	11	41
61	28	58	24	53	18	47	12
31	62	29	59	25	54	19	48
64	32	63	30	60	26	55	20

Forward Substitution: 2nd Color of CM-RCM(2)

Sliced-ELL Row-wise

SCS-a (SELL-C-σ) Row-wise

SCS-b (SELL-C-σ) Column-wise

(a) Sliced ELL: Row-Wise

```
!$omp parallel private (icol,…)
    icol= NCOLORtot
    (Operations)
!$omp do
    do (Loops:Meshes: Row-Wise)
    do (Loops:Non-Zero Off-Diag's: j= 1, 6)
        (Operations)
    enddo
    enddo
!$omp end parallel
```
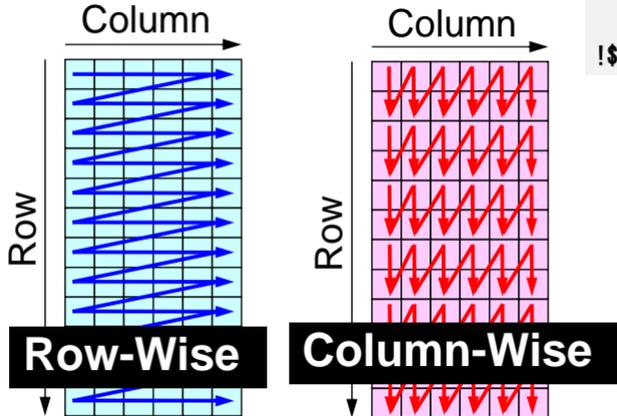
(b) SELL-8-8/Row-Wise (SCS-a)

```
!$omp parallel private (icol,…)
    icol= NCOLORtot
    (Operations)
!$omp do
    do (Loops:Threads: ip= 1, PEsmpTOT)
    do (Loops:Blocks)
!$omp simd
    do (Loops SIMD: k= 1, 8)
    do (Loops:Non-Zero Off-Diag's: j= 1, 6)
        (Operations)
    enddo
    enddo
    enddo
    enddo
!$omp end parallel
```

(c) SELL-8-8/Column-Wise (SCS-b)

```
!$omp parallel private (icol,…)
    icol= NCOLORtot
    (Operations)
!$omp do
    do (Loops:Threads: ip= 1, PEsmpTOT)
    do (Loops:Blocks)
!$omp simd
    do (Loops SIMD: k= 1, 8)
        (Operations:Substitutions-1)
    enddo
    do (Loops:Non-Zero Off-Diag's: j= 1, 6)
!$omp simd
    do (Loops SIMD: k= 1, 8)
        (Operations)
    enddo
    enddo
!$omp simd
    do (Loops SIMD: k= 1, 8)
        (Operations:Substitutions-2)
    enddo
    enddo
    enddo
!$omp end parallel
```

More Expensive !!

Column

Column

Row

Row

Row-Wise

Column-Wise

Parallel Multigrid Methods
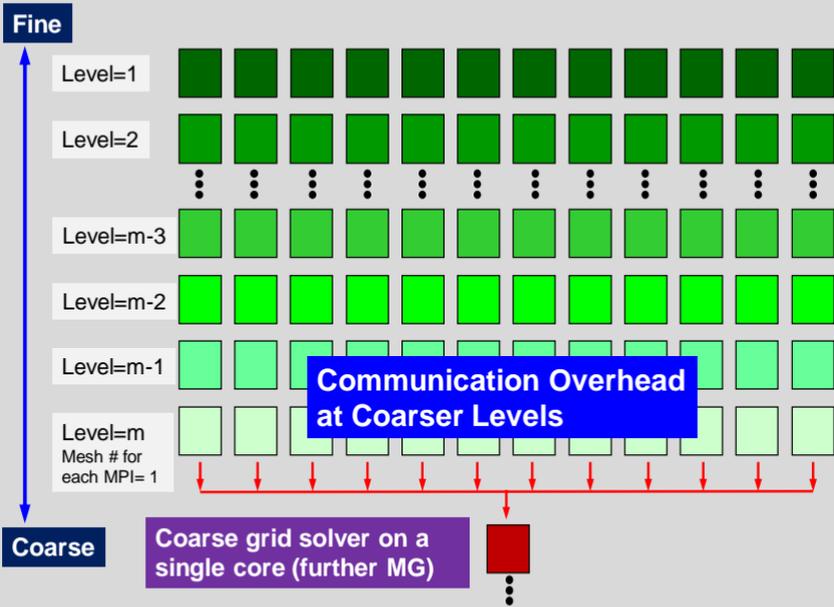
Original Approach [KN 2010]

- Fine
  - Level=1
  - Level=2
  - Level=m-3
  - Level=m-2
  - Level=m-1
  - Level=m, Mesh # for each MPI= 1
- Coarse

Communication Overhead at Coarser Levels

Coarse grid solver on a single core (further MG)
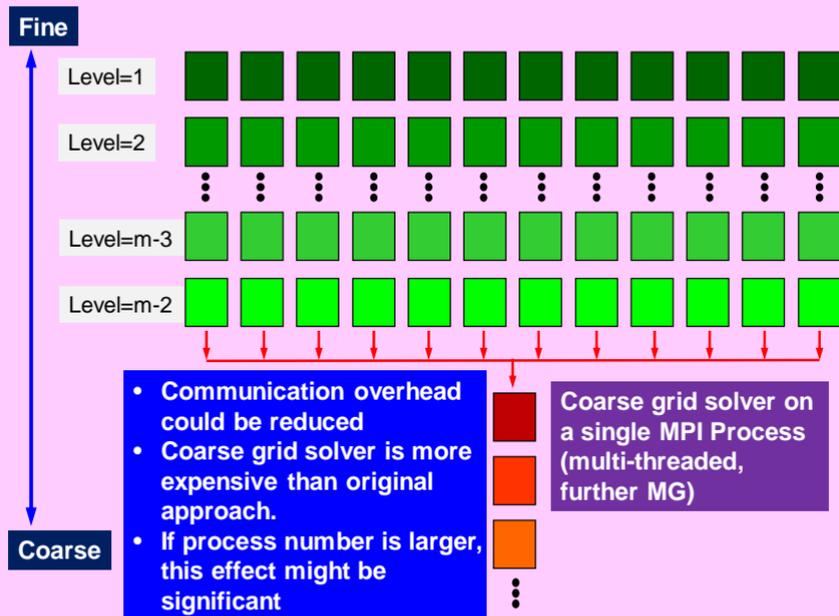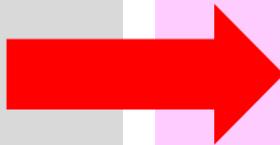
Coarse Grid Aggregation (CGA) [KN 2012]

- Fine
  - Level=1
  - Level=2
  - Level=m-3
  - Level=m-2
- Coarse

- Communication overhead could be reduced
- Coarse grid solver is more expensive than original approach.
- If process number is larger, this effect might be significant

Coarse grid solver on a single MPI Process (multi-threaded, further MG)

Summary of Configurations

- pGW3D-FVM, Weak Scaling
  - CGA（Coarse Grid Aggregation）: Single Level
- Oakforest-PACS (OFP), ~2,048 nodes
  - Flat, MC-DRAM only
  - 64-cores on each node
    - HB 4x16 (4-threads x 16-proc's), HB 8x8
  - Problem Size: 64x32x32 on each core (max: 8,589,934,592 DOF), Best for 5-Runs, $\varepsilon=10^{-12}$
- Comparison between (CRS, Sliced ELL) and SCS (SELL-C-σ, C=σ=8)
  - 64-bit (Double Precision) x 8 = 512 bit
  - SCS: Switching to Sliced-ELL if the problem size is smaller than (C (=8)) for each color/thread
  - **Sliced ELL for Coarse Grid Solver of SCS**
  - 20% improvement is expected by SCS over Sliced ELL, based on preliminary results
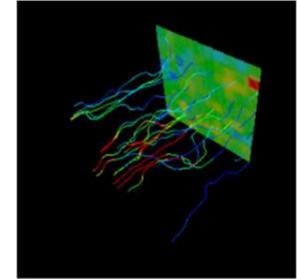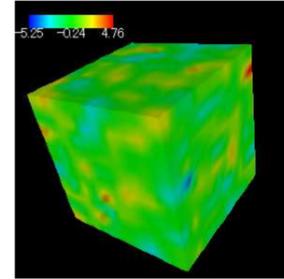
Results: Time for MGCG
8-nodes, HB 4×16, 33,554,432 DOF
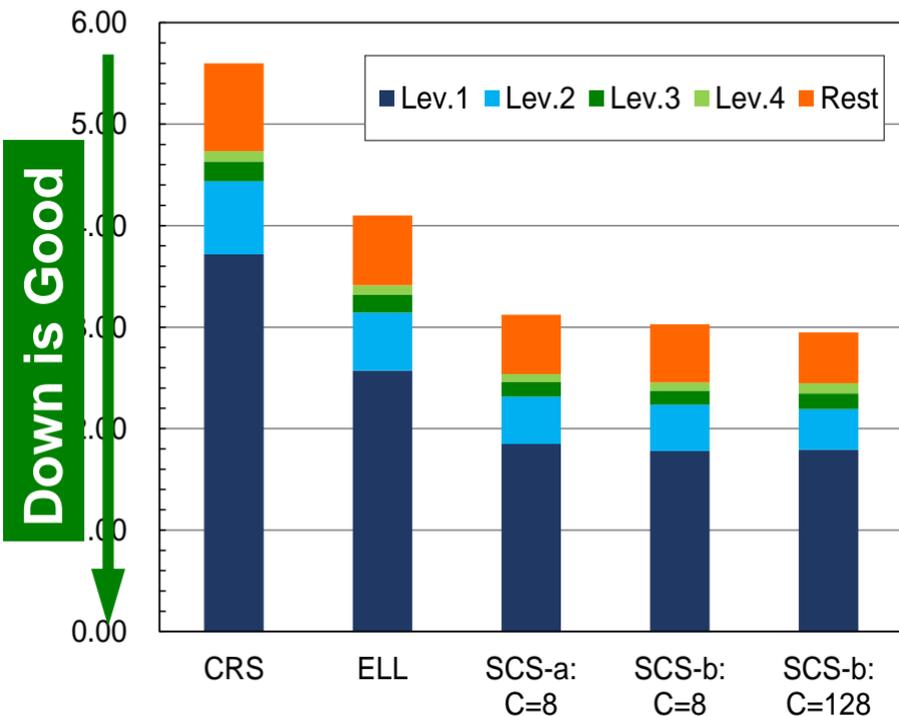
- Each Level of Smoothing
- Rest
  - Coarser Level Smoothers
  - CG except MG (SpMV etc.)
  - Communication
  - Coarse Grid Solver
- Improvement over CRS
  - Sliced ELL：36.6%
  - SCS-a：79.4%
  - SCS-b：84.9% (C=8)
  - SCS-b: 90.1% (C=128)
- Level-1 (Finest)
  - 46.6%, 101.3%, 108.7%, 107.8%
- Sliced ELL⇒SCS
  - SCS-a：31.4%
  - SCS-b：35.4% (C=8)
  - SCS-b: 39.2% (C=128)

Weak Scaling: up to 2,048-nodes
HB 4x16, Time for MGCG, Down is Good
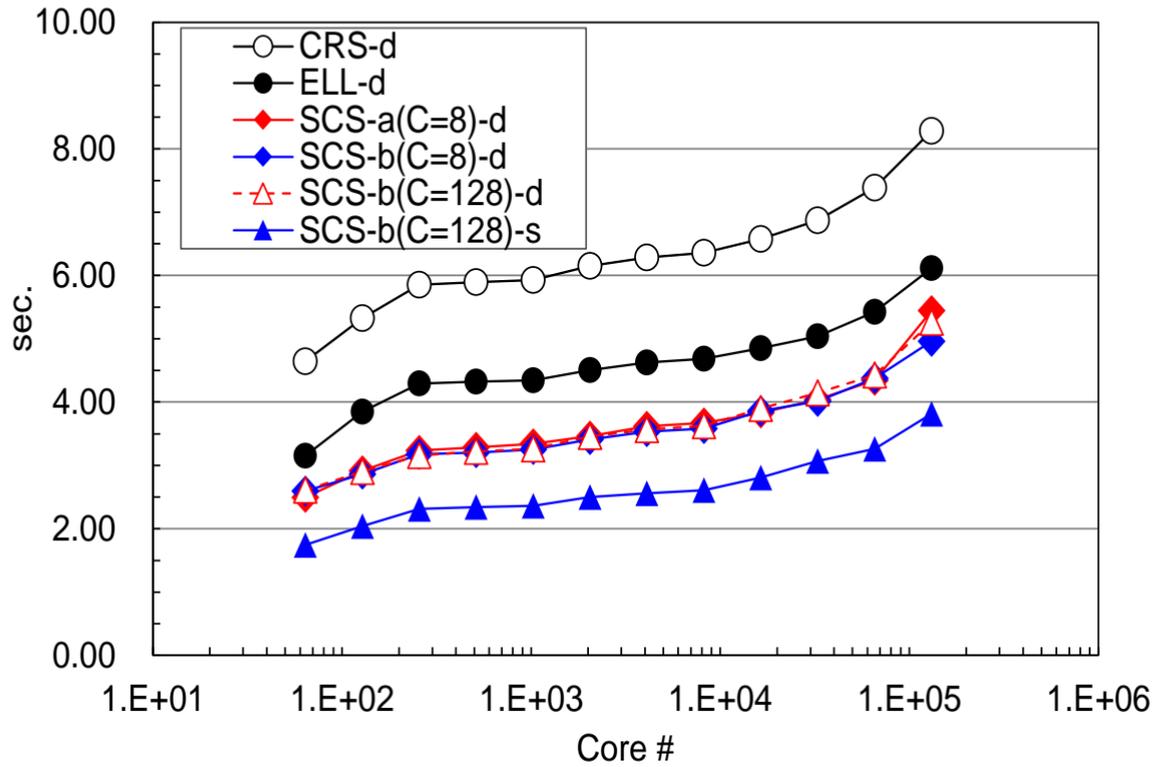
Weak Scaling: up to 2,048-nodes
SCS-b (C=σ=8), Time for MGCG, <u>Down is Good</u>

HB 4x16
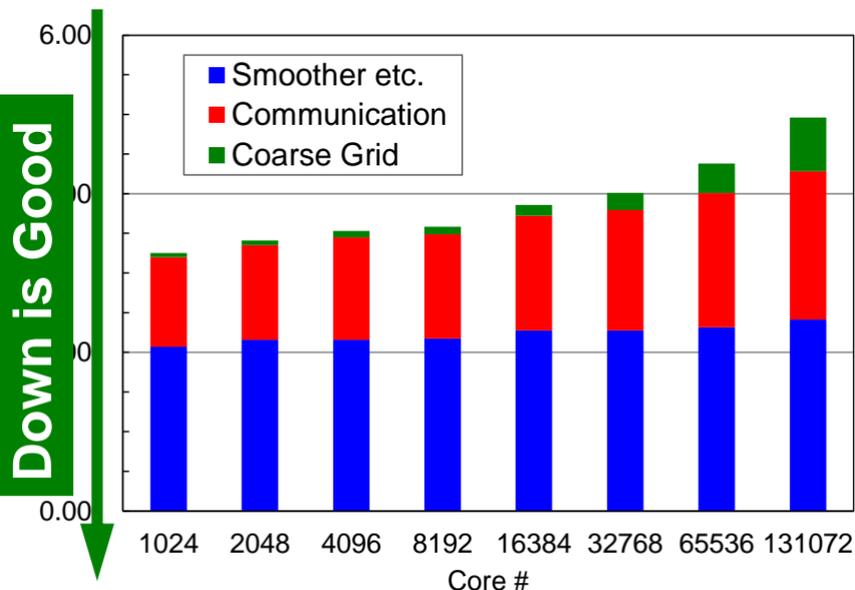
HB 8x8
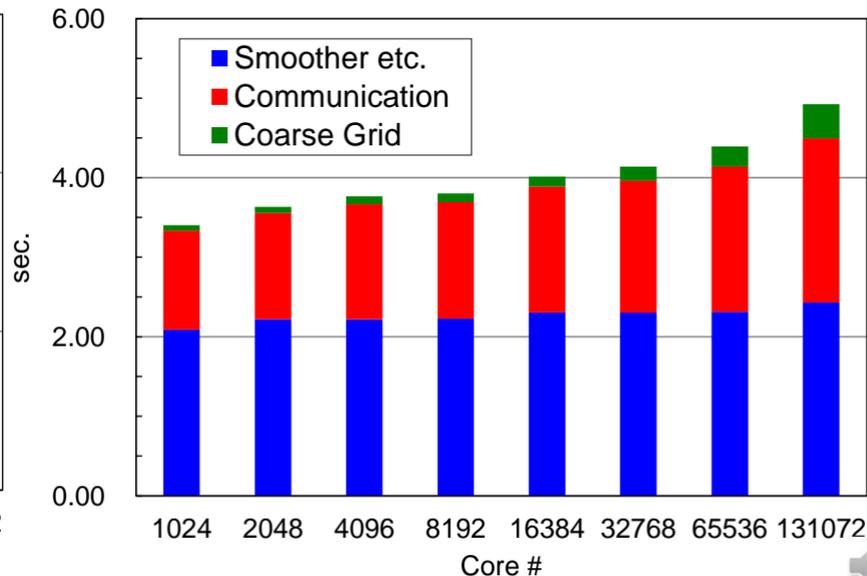
Down is Good

- Background & Overview
  - Multigtrid Methods
  - Overview of Previous/Present Works
- SELL-C-σ with Double Precision Computing
- **Computing in Double/Single Precision**
- Accuracy Verification
- Summary

Approximate Computing with Low/Adaptive/Trans Precision

- Mostly, scientific computing has been conducted using FP64 (double precision, DP)
  - Sometimes, problems can be solved by FP32 (single precision, SP) or lower precision
- **Lower precision may save time, energy and memory**
- Approximate Computing
  - Originally for image recognition etc. where accuracy is not necessarily required
  - Also applied to numerical computations
- Computations by lower precision and by mixed precision may provide results with less accuracy

P3D: Steady State 3D Heat Conduction by FVM

$$\nabla \cdot \left(\lambda \nabla \phi \right) + f = 0$$

- 7-point Stencil
- Heterogenous Material Property
  - λ_1/λ_2 is proportional to the condition number of coefficient matrices
- Coefficient Matrix
  - Sparse, SPD
- ICCG Solver
- Fortran 90 + OpenMP
- CM-RCM Reordering
- FP64 (Double), FP32 (Single), FP16 (Half) (just for preconditioning)

Ratio of FP32(SP)/FP64(DP)

Iterations● & Time△ for ICCG
λ_1/λ_2, 128^3 DOF, CRS
Ratio<1 ⇒ FP32 is faster

λ_1

λ_2

λ_1

$$\nabla \cdot (\lambda \nabla \phi) + f = 0$$
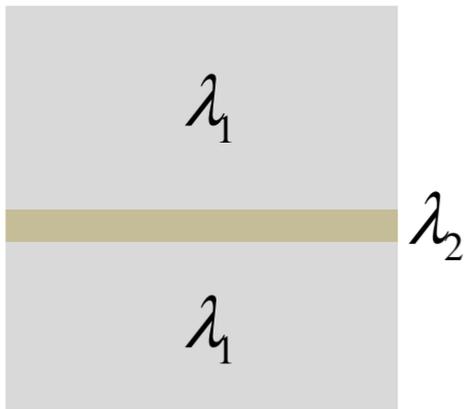
Intel Xeon BDW
1 Node: 18 cores x 2 soc's

Ratio of FP32/FP64

●Iterations △Time

Ratio of λ_1/λ_2

[KN et al. 2018]

Ratio of FP32(SP)/FP64(DP)

Iterations● & Time△ for ICCG

λ_1/λ_2, 128^3 DOF, CRS

Ratio<1 ⇒ FP32 is faster

$$\nabla \cdot \left(\lambda \nabla \phi \right) + f = 0$$

λ_1

λ_2

λ_1

Intel Xeon BDW
1 Node: 18 cores x 2 soc's

●Iterations △Time

FP32: Slower

FP32: faster

Ratio of FP32/FP64

Ratio of λ_1/λ_2

[KN et al. 2018]

Ratio of FP32(SP)/FP64(DP)

Iterations ● & Time△ for ICCG
λ_1/λ_2, 128^3 DOF, CRS
Ratio<1 ⇒ FP32 is faster

$$\nabla \cdot (\lambda \nabla \phi) + f = 0$$

Intel Xeon BDW
1 Node: 18 cores x 2 soc's

●Iterations △Time

+20% iterations
by FP32

-40~-45% Time by FP32

Ratio of FP32/FP64

Ratio of λ_1/λ_2

[KN et al. 2018]

Ratio of FP32(SP)/FP64(DP)

Iterations● & Time△ for ICCG

λ_1/λ_2, 128^3 DOF, CRS

Ratio<1 ⇒ FP32 is faster

$$\nabla \cdot \left(\lambda \nabla \phi \right) + f = 0$$

●Iterations △Time

+20% iterations by FP32

-40~-45% Time by FP32

Ratio of FP32/FP64

Ratio of λ_1/λ_2

1.E+00   1.E+01   1.E+02   1.E+03   1.E+04   1.E+05   1.E+06

[KN et al. 2018]

Results on Intel Xeon BDW $\lambda_1 = \lambda_2$ [Sakamoto et al. 2020]
N=128³, ■: CPU, ■: Memory, ●:Time
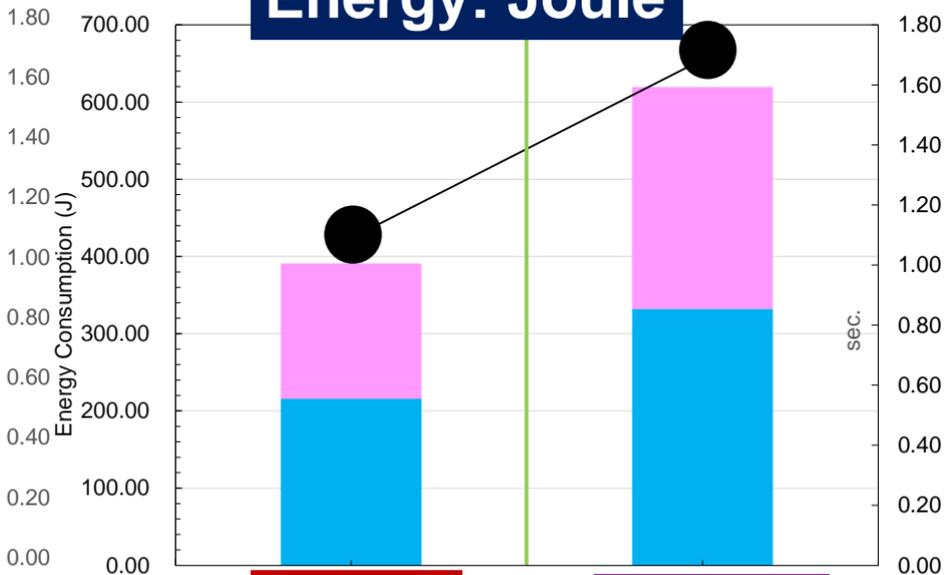
Target Platforms

- Oakforest-PACS (OFP)
  - Intel Xeon Phi (Knights Landing, KNL), Fujitsu
  - IHK/McKernel
  - 8,208 nodes, 25+PF, 22th in TOP 500 (Nov.2020)
  - Operated by JCAHPC (U.Tsukuba & U.Tokyo)
- Oakbridge-CX (OBCX)
  - Intel Platinum 8280 (Cascade Lake, CLX), Fujitsu
  - 1,368 nodes (2,736 sockets), 6.61 PF, 69th in TOP 500 (Nov.2020)

Overview of Each Node (OFP & OBCX)

System	Oakforest-PACS (OFP)	Oakbridge-CX (OBCX)
Name in this Paper	OFP	OBCX
Architecture	Intel Xeon Phi 7250 (Knights Landing, KNL)	Intel Xeon Platinum 8280 (Cascade Lake, CLX)
Frequency (GHz)	1.40	2.70
Core #/CPU (socket)	68	28
CPU (socket) # per node	1	2
Peak Performance (GFLOPS) per node	**3,046.4**	**4,838.4**
Memory Size (GB) per node	MCDRAM: 16 DDR4: 96	192
Memory Bandwidth/Socket (GB/sec, STREAM Triad)	**MCDRAM: 490 DDR4: 84.5**	**202.0**
Peak Performance per Core (GFLOPS)	44.8	86.4
Memory Bandwidth per Core (GB/sec., STREAM Triad)	MCDRAM: 7.21 DDR4: 1.24	3.61

Time for MGCG: SCS-b (C=σ) Normalized by SELL-8-8 (DP)

**8-nodes
OFP(HB 4x16),
OBCX (6x8)**

All variables, vectors and matrices are stored in double precision for FP64, and all of them are stored in single precision for FP32
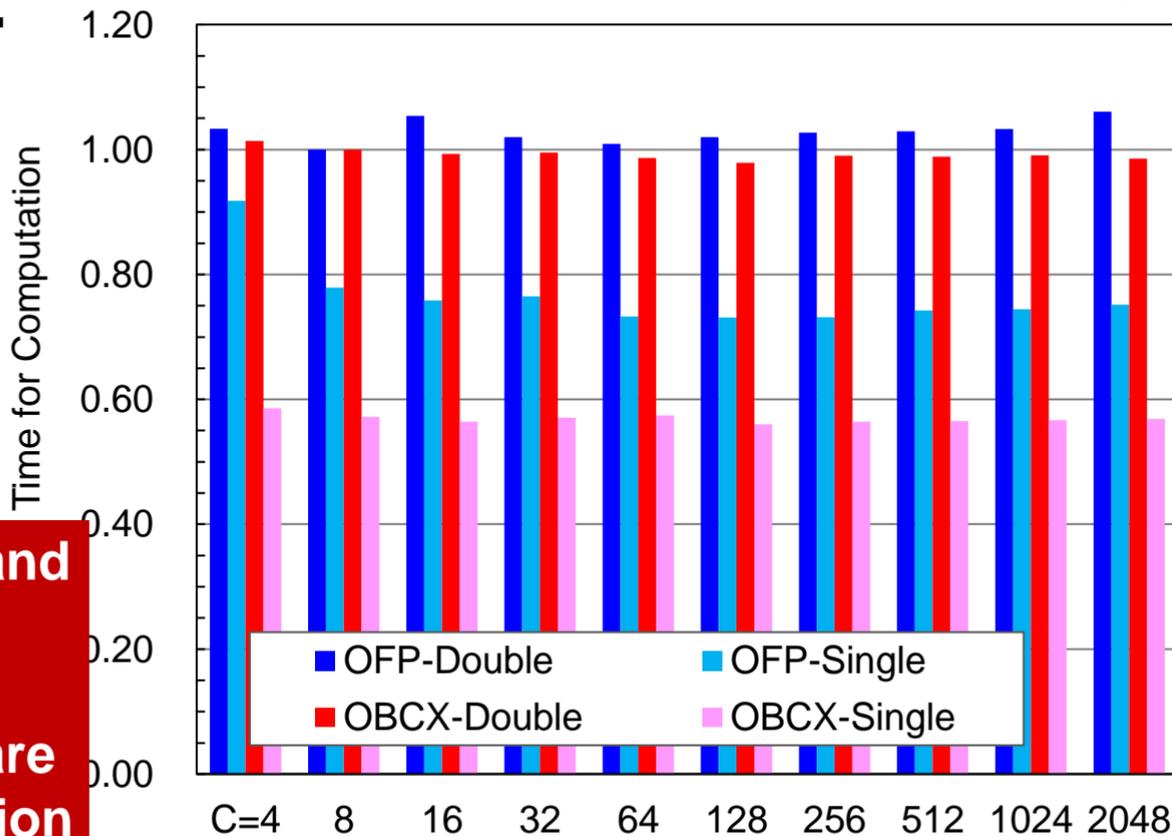
Time for MGCG: SCS-b (C=σ) Normalized by SELL-8-8 (DP)

8-nodes
OFP(HB 4x16),
OBCX (8x8)

Optimum Parameters
- DP: C=σ=   8
- SP: C=σ=128

Down is Good

Normalized Time for Computation

	OFP-Double	OFP-Single
	OBCX-Double	OBCX-Single

C=4 8 16 32 64 128 256 512 1024 2048
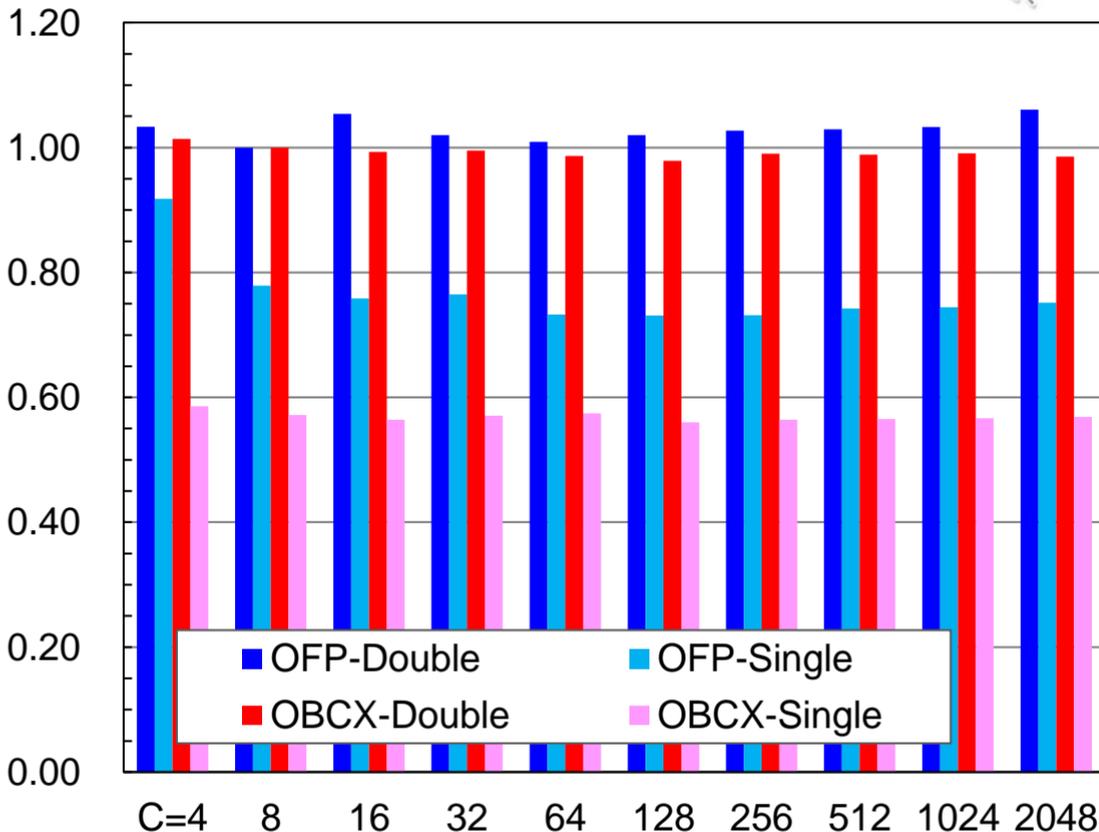
DP⇒SP

- Iteration number does not change
- <0.1% relative error
- (SP/DP) time ratio
  - 0.70 for OFP
  - 0.55-0.60 for OBCX

Normalized Time for Computation

Legend	
■ OFP-Double	■ OFP-Single
■ OBCX-Double	■ OBCX-Single

C=4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048
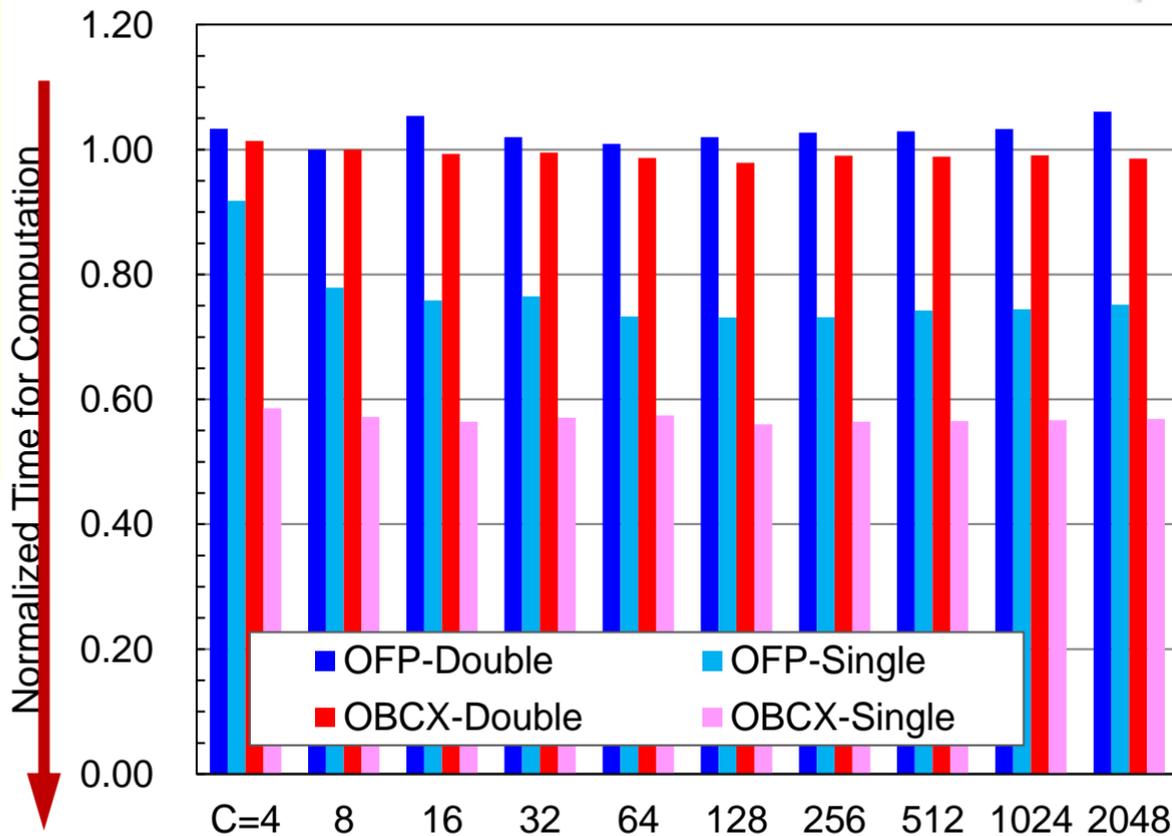
Down is Good

Results: Time for MGCG
8-nodes, OFP (HB 4 × 16)

FP64, Double

FP32, Single

Down is Good
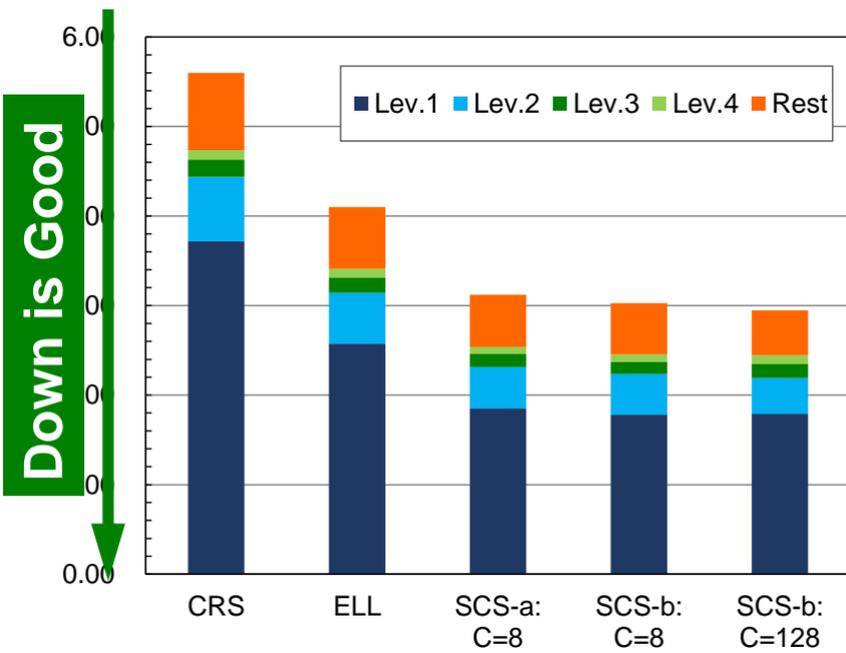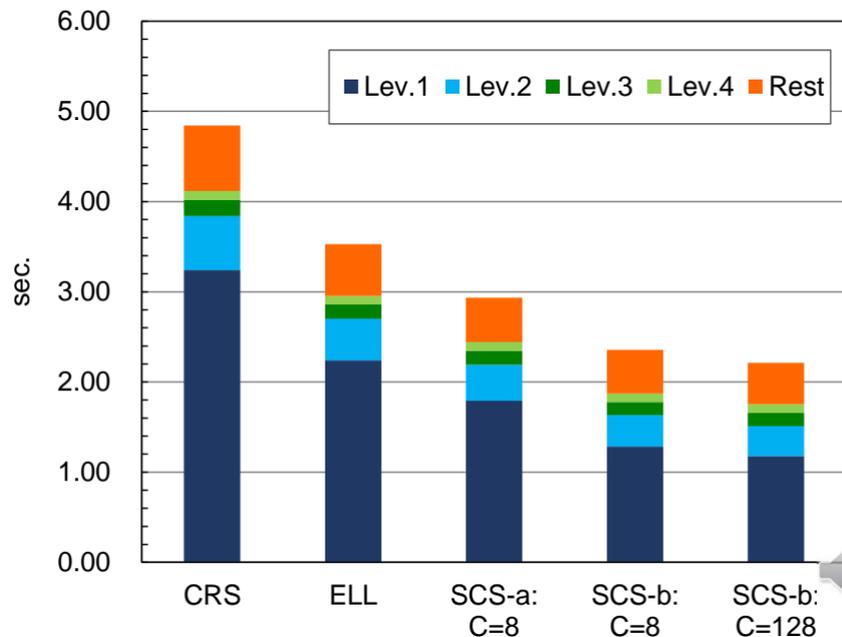
Time for MGCG, 8-nodes, OFP (HB 4×16)
Improvement over CRS (MGCG solver, Level-1 of Smoother)

	Double Precision (FP64)	Single Precision (FP32)
Sliced ELL	36.6%, 46.6%	15.6%, 44.7%
SCS-a (C=σ=8)	79.4%, 101.3%	58.7%, 81.0%
SCS-b (C=σ=8)	84.9%, 108.7%	90.9 %, 152.2%
SCS-b (C=σ=128)	90.1%, 107.8%	**137.5%, 174.9%**

Results: Time for MGCG
8-nodes, OBCX (HB 6×8), Down is Good

FP64, Double
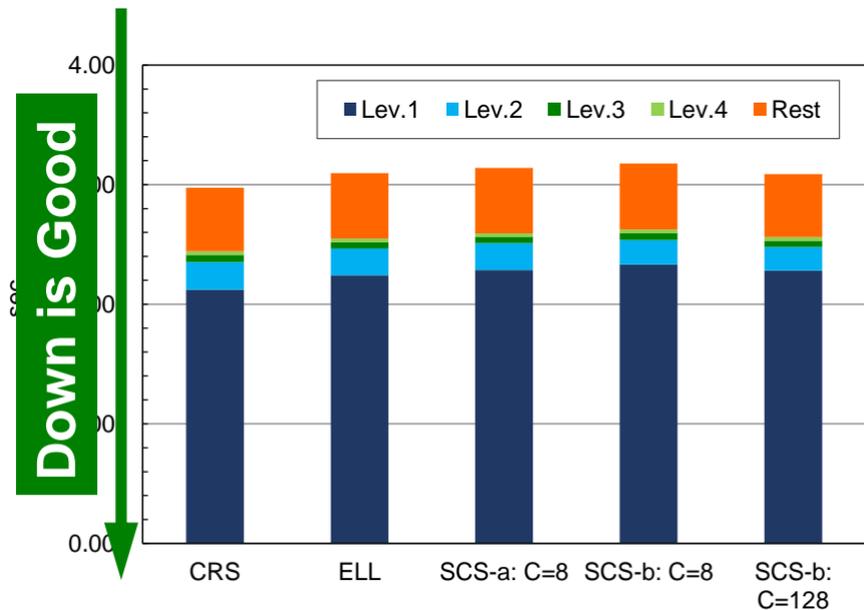
FP32, Single
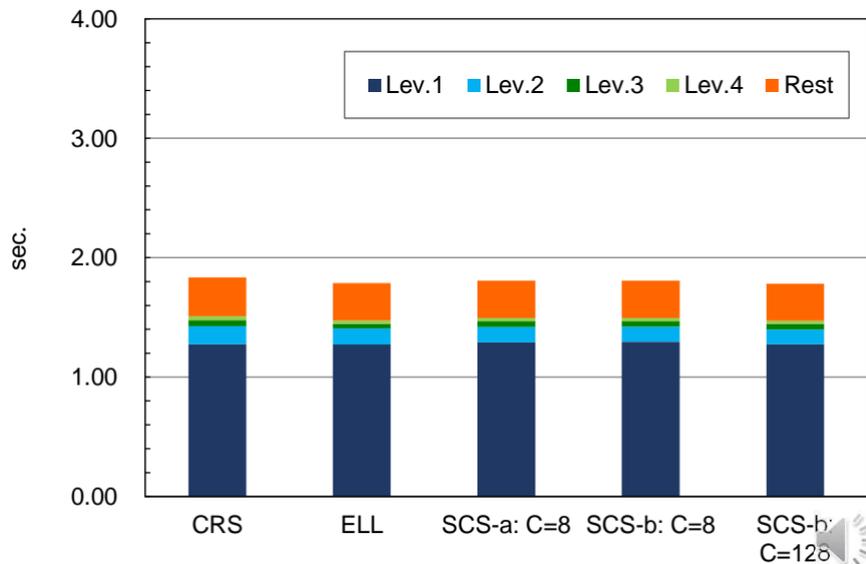
Weak Scaling: up to 2,048-nodes of OFP (DP: C=σ=8, SP: C=σ=128)
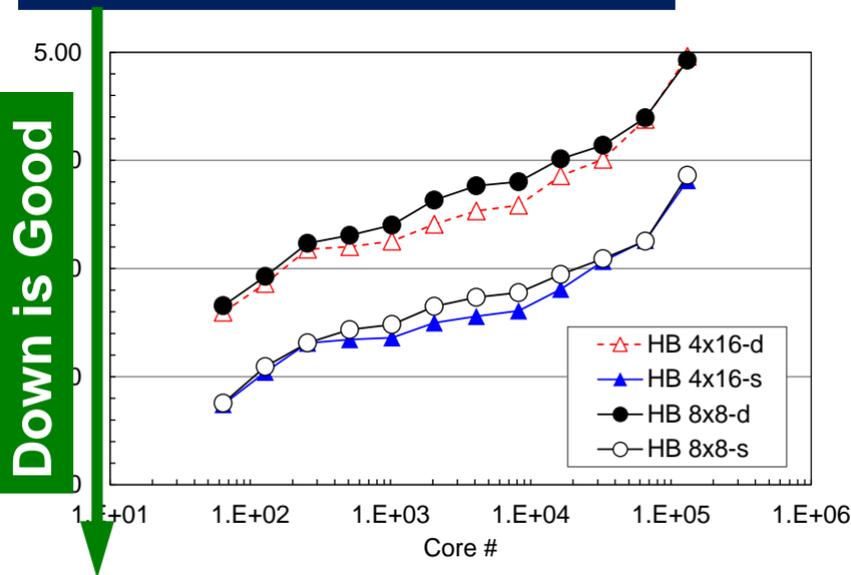
Time for MGCG: SCS-b
DP: C=σ=8, SP: C=σ=128

Down is Good

Legend:
- HB 4x16-d
- HB 4x16-s
- HB 8x8-d
- HB 8x8-s

X-axis: Core #

Weak Scaling: up to 2,048-nodes of OFP (DP: C=σ=8, SP: C=σ=128)

Time for MGCG: SCS-b
DP: C=σ=8, SP: C=σ=128
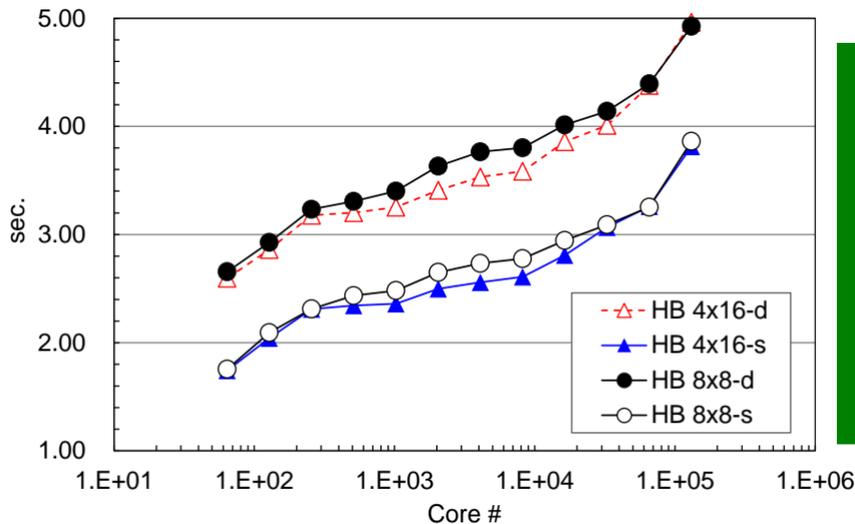
Improvement over CRS

Up is Good
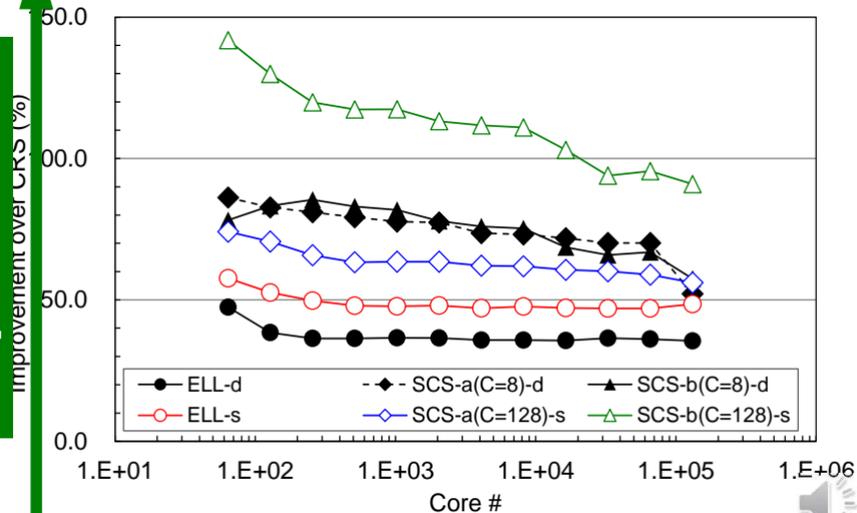
Weak Scaling: up to 2,048-nodes of OFP, HB 4x16, Time for MGCG, <u>Down is Good</u>

FP64: C=σ=8
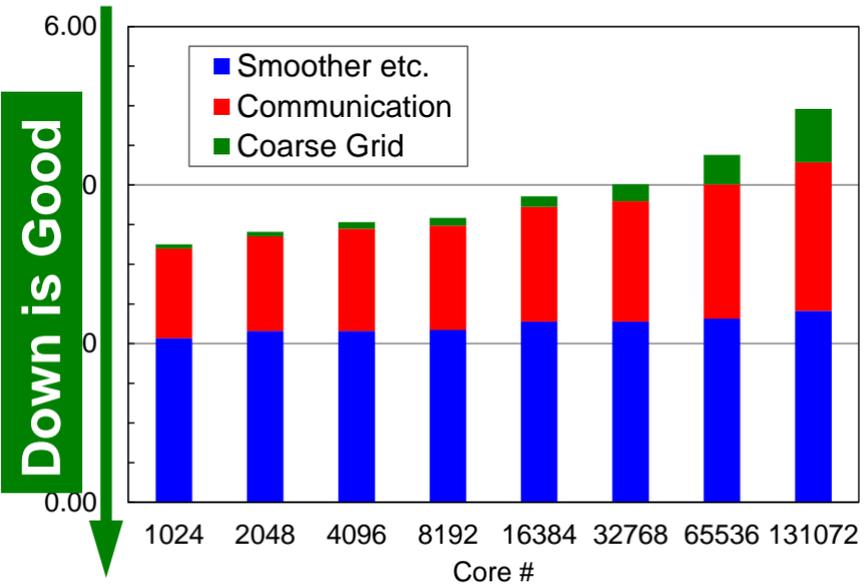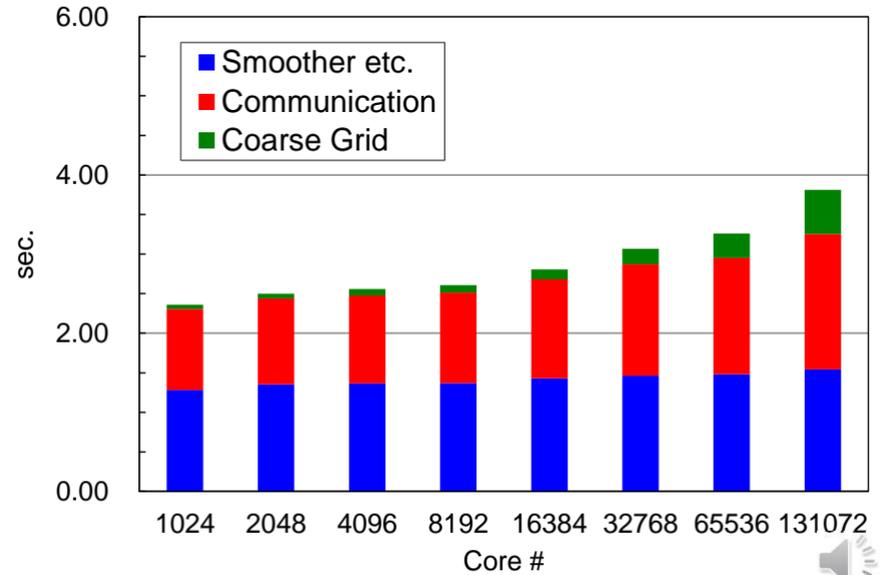
FP32: C=σ=128

- Background & Overview
  - Multigtrid Methods
  - Overview of Previous/Present Works
- SELL-C-σ with Double Precision Computing
- Computing in Double/Single Precision
- **Accuracy Verification**
- Summary

Approximate Computing with Low/Adaptive/Trans Precision

- Accuracy verification is important, especially for computation with lower/mixed precision.

- A lot of methods for accuracy verification have been developed for problems with dense matrices
  - But very few examples for sparse matrices & H-matrices

- Generally speaking, processes for accuracy verification is very expensive
  - Sophisticated Method needed
  - Automatic Selection of Optimum Precision by Technology of AT (Auto Tuning)

- Accuracy Verification of Sparse Linear Solvers [Ogita, Nakajima 2019]

New Algorithm for Verification
[Ogita, Rump, Oishi 2005] [Ogita, Nakajima 2019]

1. Solve a discretized linear system $Ax = b$.

2. Solve a linear system $Ay = e$.

3. Verify M-property of A using $\hat{y}$. $(\hat{y} > 0 \Rightarrow A\hat{y} > 0)$

4. Compute $r = b - A\hat{x}$ with an error bound.

   ➤ $\hat{r}$: a computed residual, e_r: an error bound of $\hat{r}$

5. Solve a linear system $Az = \hat{r}$.

6. Compute an error bound using

$$\|x - \hat{x}\|_\infty \leq \|\hat{z}\|_\infty + \frac{\|\hat{y}\|_\infty(\|\hat{r} - A\hat{z}\|_\infty + \|e_r\|_\infty)}{1 - \|e - A\hat{y}\|_\infty}.$$

Results: New Alg. for Verification (128^3)
[Ogita, Rump, Oishi 2005] [Ogita, Nakajima 2019]

$$\phi = 0 \,@\, z = z_{max}$$

NZ

λ_1

λ_2

λ_1

$$\nabla \cdot (\lambda \nabla \phi) + f = 0$$

better

- ● Improved error bound
- ■ Maximum relative error
- ▲ Relative residual norm

10^{-5}

10^{-4}

lambda1 / lambda2

Previous bound

New error bound

$$\max_{1 \le i \le n} \left| \frac{x_i - \hat{x}_i}{x_i} \right|$$

Residual norm

$$\frac{\| b - A\hat{x} \|_2}{\| b \|_2}$$
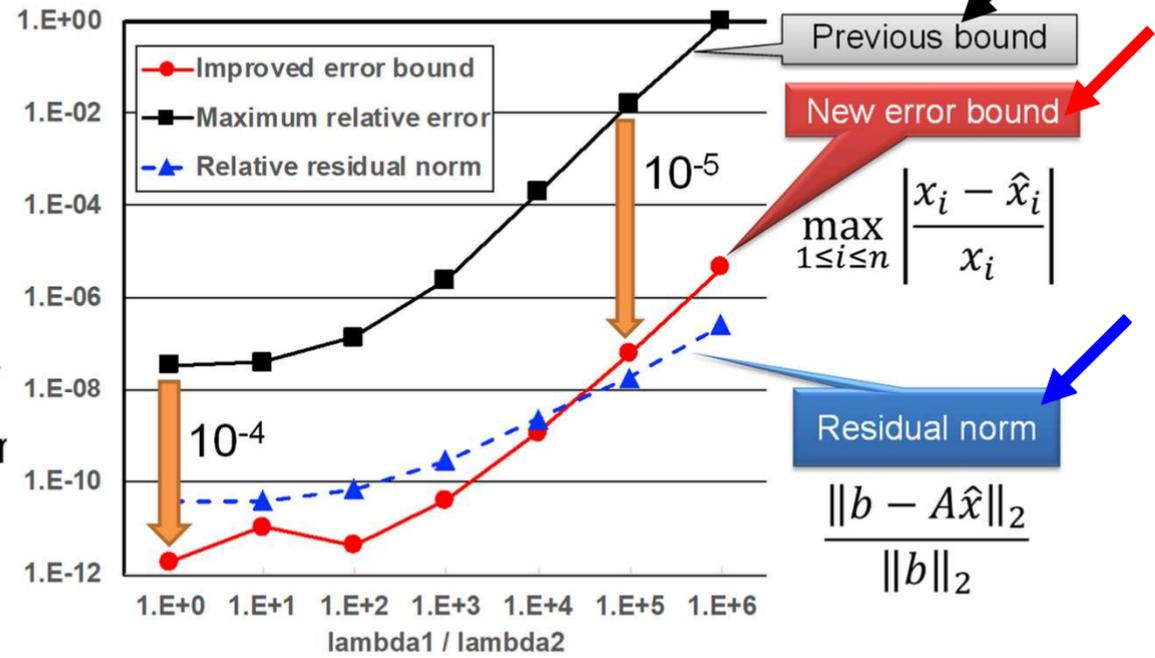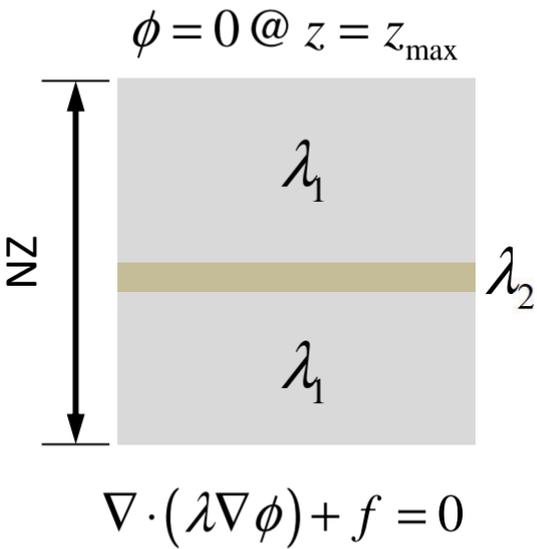
Computed error bounds are significantly improved!

Results on OBCX (Intel Xeon CXL) (1/2)

FP64, Accuracy Verification ■ takes 10% longer

$\lambda_1/\lambda_2 = 10^0 \sim 10^6$, CRS, N=$128^3$

$\phi = 0 @ z = z_{\max}$

λ_1

λ_2

λ_1

NZ

$\nabla \cdot (\lambda \nabla \phi) + f = 0$

■ : P3D Solve
■ : Acc. Veri.
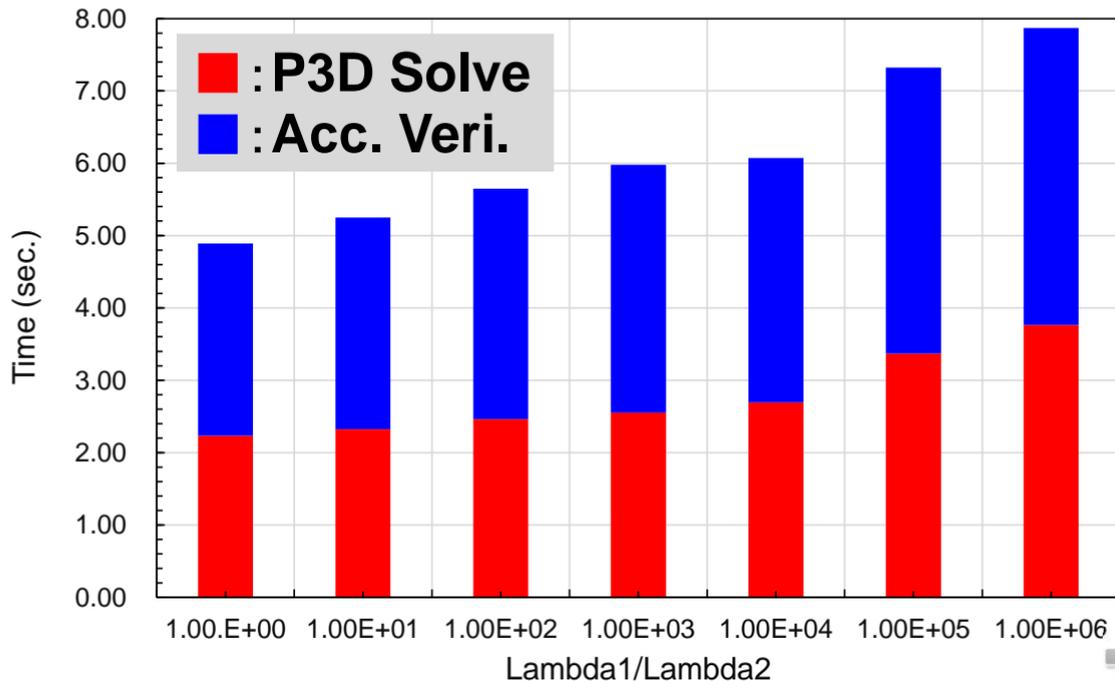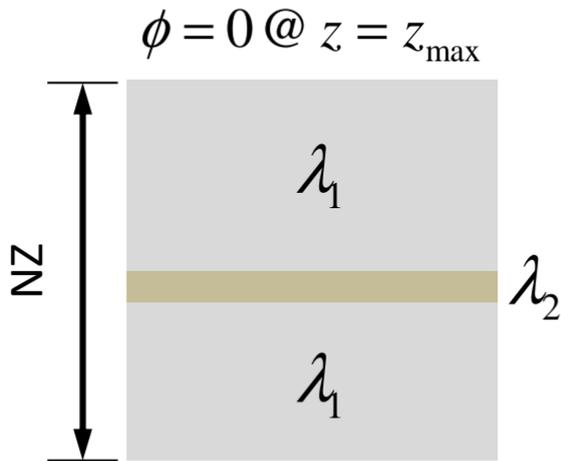
Results on OBCX (2/2)
Accuracy verification for FP32 has failed, if $\lambda_1 / \lambda_2 \geqq 10^4$

$$\|x - \hat{x}\|_\infty \leq \|\hat{z}\|_\infty + \frac{\|\hat{y}\|_\infty \|b - A(\hat{x} + \hat{z})\|_\infty}{1 - \|e - A\hat{y}\|_\infty}$$
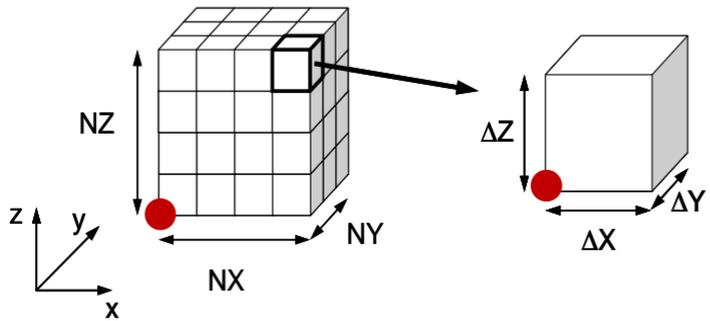
- $\bigcirc$ Relative Error between FP32 & FP64 at $\bullet$
- $\blacktriangle$ Max. Relative Error Bound for FP32 obtained from Accuracy Verification
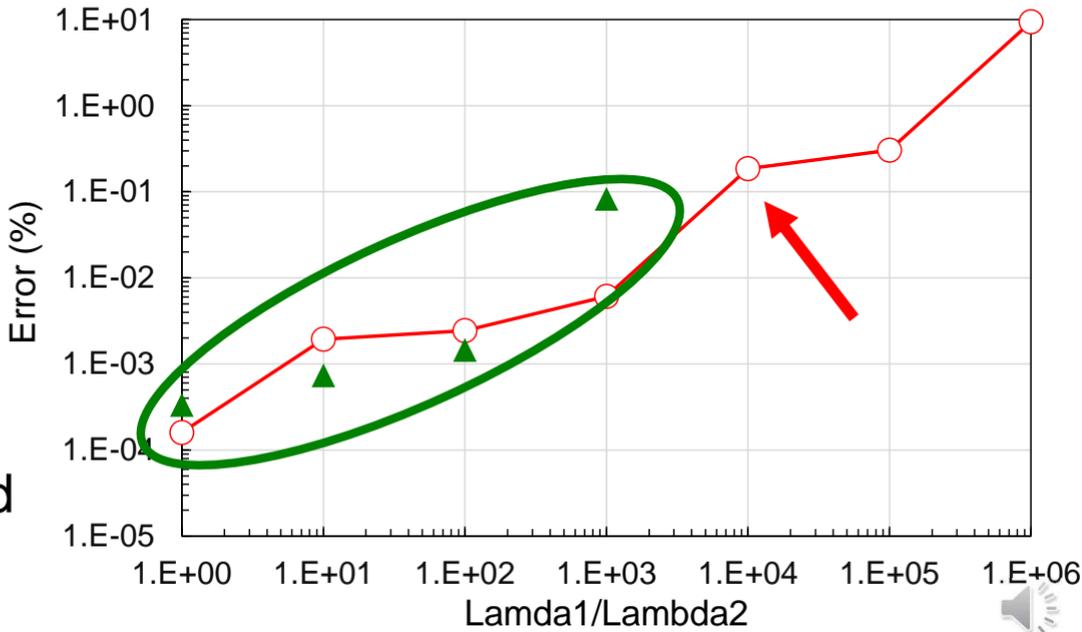
Accuracy Verification in pGW3D-FVM

- 1st Application of the New Verification Method [Ogita, Nakajima 2019] to Distributed Parallel Computing
- Two Cases on OFP (HB 8x8)
  - Small: 128 meshes, 1-node
  - Large: 1,024x1,024x512 meshes, 128 nodes
- DP only (SP failed in Accuracy Verification)
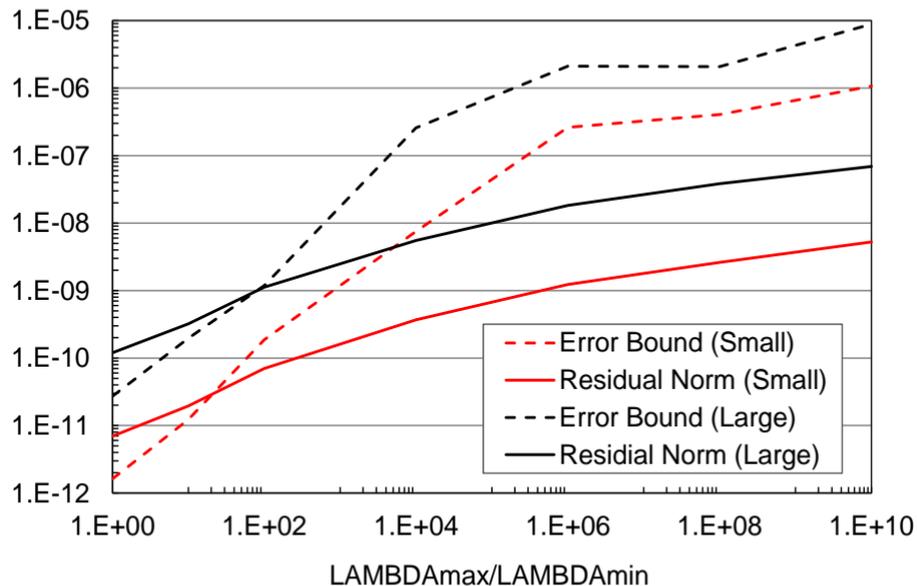
Legend:
- Error Bound (Small)
- Residual Norm (Small)
- Error Bound (Large)
- Residial Norm (Large)

X-axis: LAMBDAmax/LAMBDAmin

- Background & Overview
  - Multigtrid Methods
  - Overview of Previous/Present Works
- SELL-C-σ with Double Precision Computing
- Computing in Double/Single Precision
- Accuracy Verification
- **Summary**

Summary

- SELL-C-σ with double/single precision computing in the MGCG solver using up to 2,048 nodes of Intel Xeon Phi.
- Improvement of the performance by SELL-C-σ over the sliced ELL
  – 35+% for DP, 45+% for SP on OFP.
  – The effect of SELL-C-σ for computing with SP is very significant
- Accuracy Verification (Preliminary)
- The first example of SELL-C-σ applied to forward/backward substitutions in ILU-type smoother of multigrid solver with double/single precision computing.
- Future Works
  – SELL-C-σ for Coarse Grid Solver
  – Improvement of Accuracy Verification