Numerical Library with High-Performance/Adaptive-Precision/High-Reliability

Extension of ppOpen-HPC/ESSEX-II towards the Post Moore Era

Kengo Nakajima[1,3] Takeshi Ogita[2)]
Ryuichi Sakamoto[1)]

1) Information Technology Center, The University of Tokyo
2) Tokyo Woman's Christian University 3) RIKEN R-CCS

4th SPPEXA Workshop on Parallel Programing Models - Productivity & Applications for Exascale and Beyond
University of Versailles, France, March 21, 2019

Acknowledgements

- Sponsors
  - ✓ CREST-JST, Japan
  - ✓ SPPEXA-DFG, Germany
  - ✓ JHPCN, Japan
- Collaborators, Colleagues
  - ✓ Takeshi Iwashita (Hokkaido U.)
  - ✓ Takahiro Katagiri (Nagoya U.)
  - ✓ Takashi Shimokawabe (ITC/U.Tokyo)
  - ✓ Hisashi Yashiro (RIKEN R-CCS)
  - ✓ Hiroya Matsuba (RIKEN R-CCS)
  - ✓ Hiromichi Nagao (ERI/U.Tokyo)
  - ✓ Takeshi Ogita (TWCU)
  - ✓ Ryuichi Sakamoto (ITC/U.Tokyo)
  - ✓ Toshihiro Hanawa (ITC/U.Tokyo)
  - ✓ Akihiro Ida (ITC/U.Tokyo)
  - ✓ Tetsuya Hoshino (ITC/U.Tokyo)
  - ✓ Masatoshi Kawai (RIKEN R-CCS)
  - ✓ Takashi Furumura (ERI/U.Tokyo)
  - ✓ Hajime Yamamoto (Taisei)
  - ✓ Gerhard Wellein (Erlangen)
  - ✓ Achim Basermann (DLR)
  - ✓ Osni Marques (LBNL)
  - ✓ Weichung Wang (NTU, Taiwan)

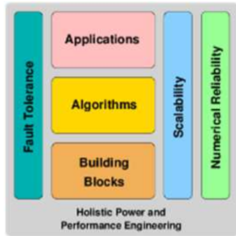

Post-Peta CREST ⇒ SPPEXA

- ppOpen-HPC (FY.2011-2015) (Leading PI)
  - ✓ Open source infrastructure for development and execution of large-scale scientific applications on post-peta-scale supercomputers with automatic tuning
  - ✓ Application Framework with AT
  - ✓ https://github.com/Post-Peta-Crest/ppOpenHPC
- ESSEX-II (FY.2016-2018) (Co-PI)
  - ✓ Preconditioned Iterative Solver for Eigenvalue Problems in Quantum Science

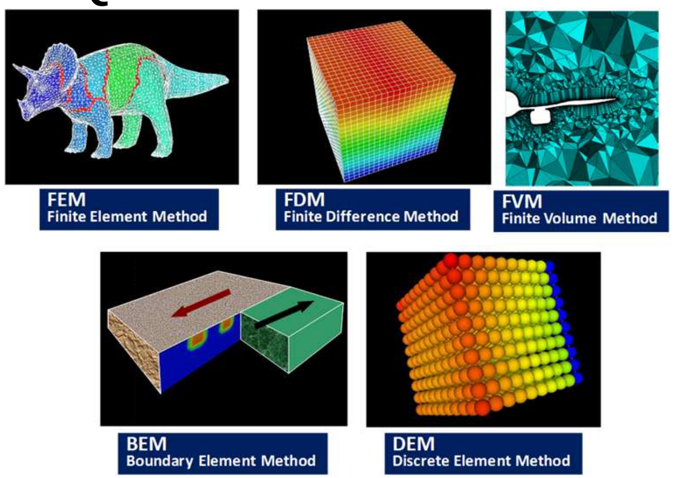

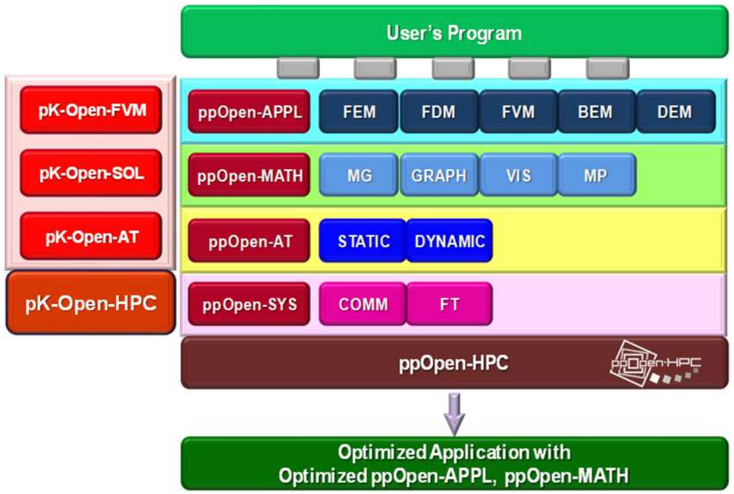

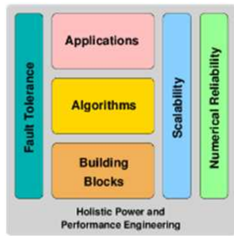

Featured Developments

- ppOpen-AT: AT Language for Loop Optimization
  - Focusing on Optimum Memory Access
- HACApK library for H-matrix comp. in ppOpen-APPL/BEM (OpenMP/MPI Hybrid Version)
  - First Open Source Library by OpenMP/MPI Hybrid
- **ppOpen-MATH/MP (Coupler for Multiphysics Simulations, Loose Coupling of FEM & FDM)**
- Sparse Linear Solvers

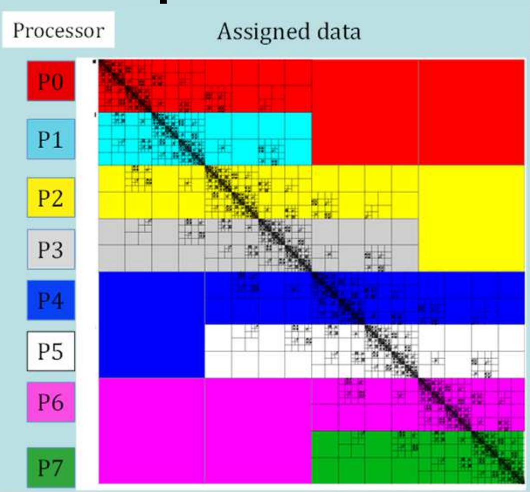

Assign a^c to P_k if $S(a^c{}_{ij})$ in $R(P_k)$

A : Whole matrix
A_{IJ} : The entry in the I-th row and the J-th column of A
a^c: A small submatrix of A
$a^c{}_{ij}$: The entry in the i-th row and the j-th column of a^c
N : Number of rows of A
n : Number of processors
P_k : The k-th processor
$l_k : 1 = l_0 < \cdots < l_k < \cdots < l_{n+1} = N + 1$
$R : R(P_k) = \{i \mid l_k \leq i < l_{k+1}\}$
$S : S(a^c{}_{ij}) = I$ when $a^c{}_{ij} = A_{IJ}$

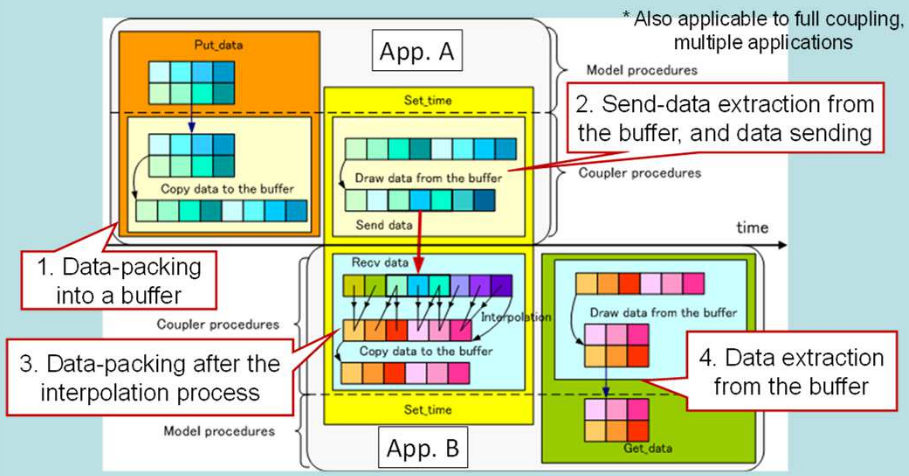

Computing in the Exascale/Post Moore Era

- Power Consumption is the Most Important Issue in the Post Moore Era
  - It is already important now.
  - Memory performance in the Post Moore Era is relatively better than now, but data movement should be reduced from the view point of energy consumption.
- Integration of (Simulation+Data+Learning) (A21 DOE)
- Quantum Computing, FPGA ?: "Partial" Solution
  - Could be a solution in certain applications (e.g. searching, graph, data clustering etc.)
  - Contributions to (S+**D+L**)
- **How to save Energy for Sustainability ?**
  - **(1) Approximate Computing by Low/Adaptive Precision**
  - **(2) Reduction of Computations: Data Driven Approach**

Approximate Computing with Low/Adaptive/Trans Precision

- **Lower Precision: Save Time & Energy & Memory**
- Approximate Computing: originally for image recognition etc.
  - **Approach for Numerical Computations**
    - **SIAM PP18 Sessions, ICS-HPC 2018 Workshop**
  - OPRECOMP: Open transPREcision COMPuting (Horizon 2020)
- **Computations with Low Precision**
- **Mixed Precision Approach (FP16-32-64-128)**
- **Iterative Refinement**
  - **such computations may provide results with less accuracy**

JHPCN

- https://jhpcn-kyoten.itc.u-tokyo.ac.jp/en/
- **The Joint Usage/Research Center for Interdisciplinary Large-scale Information Infrastructures (JHPCN)** is made up of 8 centers of National University's equipped with supercomputers.
  - Proposal-Based, Renewed Every Year, Computational Resources Awarded (e.g. Oakforest-PACS with KNL + Tsubame 3.0 with NVIDIA P100)
- **Numerical Library with High-Performance/Adaptive-Precision/High-Reliability**
  - Staring from April 2018, as a part of JHPCN Project in Japan (Preliminary Works in FY.2018)
  - 20+ Members from 13 Institutions (Japan, Germany)
    - P.I.: Kengo Nakajima (U.Tokyo)
    - Gerhard Wellein (Erlangen), Achim Basermann (DLR)

北海道大学 HOKKAIDO UNIVERSITY
東京大学 THE UNIVERSITY OF TOKYO
東京工業大学 Tokyo Institute of Technology
Tokyo Woman's Christian University 東京女子大学
芝浦工業大學 SHIBAURA INSTITUTE OF TECHNOLOGY Established 1927
名古屋大学
九州大学 KYUSHU UNIVERSITY
RIST
RIKEN
R-CCS
FAU FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG
intel
NVIDIA
DLR Deutsches Zentrum für Luft- und Raumfahrt German Aerospace Center

Numerical Library with High-Performance/Adaptive-Precision/High-Reliability

Extension of ppOpen-HPC towards the Post Moore Era

- Lower/Adaptive Precision + Accuracy Verification
  - Iterative Refinement, Mixed Precision Computation etc.
  - Verification: Collaboration with “Pure” Applied Mathematicians
- Automatic Tuning (AT): Selection of the optimum precision, which minimizes computation time and power consumption under certain target accuracy
  - implemented to “ppOpen-HPC”.
- Preconditioned Iterative Solvers for Practical Problems with Ill-Conditioned Matrices with Adaptive Precision
  - FP16-32-64-128

Results: λ_1/λ_2 ~ Condition Number
Ratio of Iterations & Computation Time
Single/Double: Down is Good

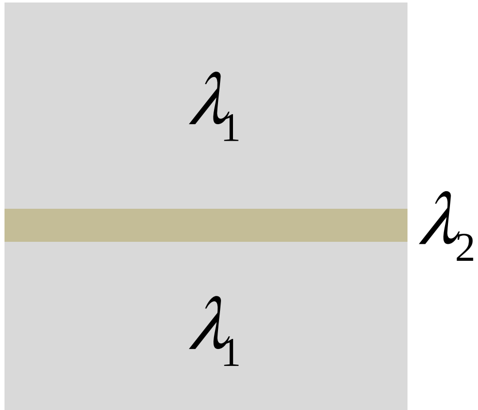

$$\nabla \cdot (\lambda \nabla \phi) + f = 0$$

Intel Xeon BDW
Single Node:
18cores x 2soc's
(Reedbush-U)

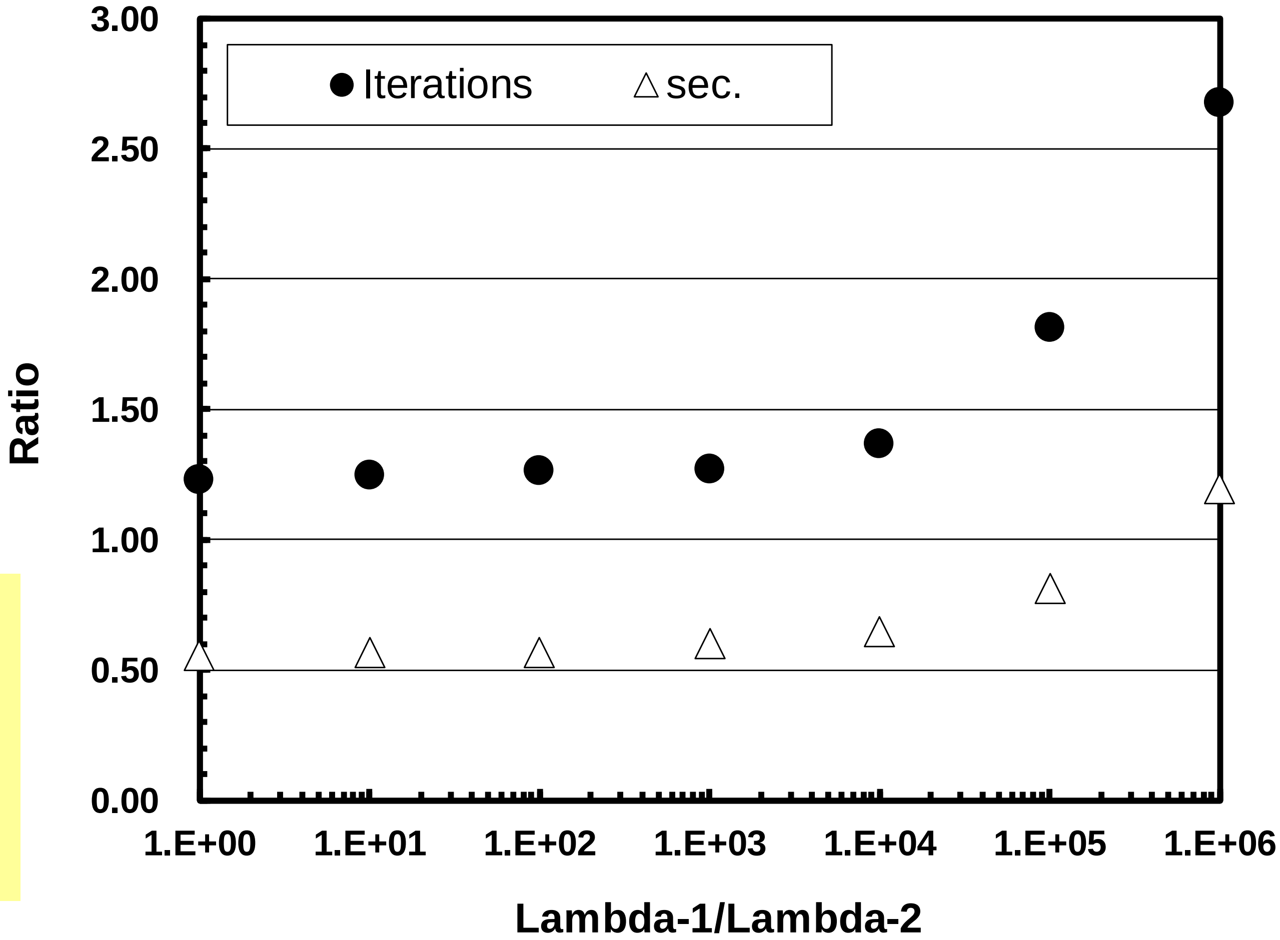

Results: λ_1/λ_2~ Condition Number
Ratio of Iterations & Computation Time
Single/Double: Down is Good

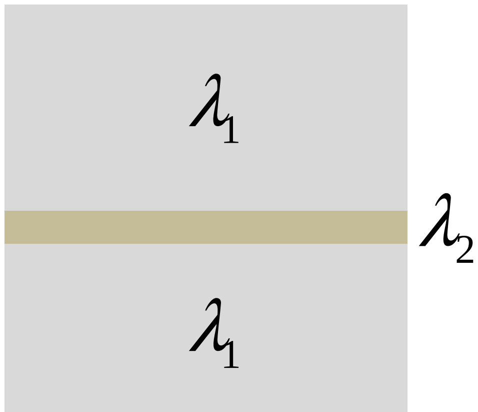

$$\nabla \cdot (\lambda \nabla \phi) + f = 0$$

Intel Xeon BDW
Single Node:
18cores x 2soc's
(Reedbush-U)

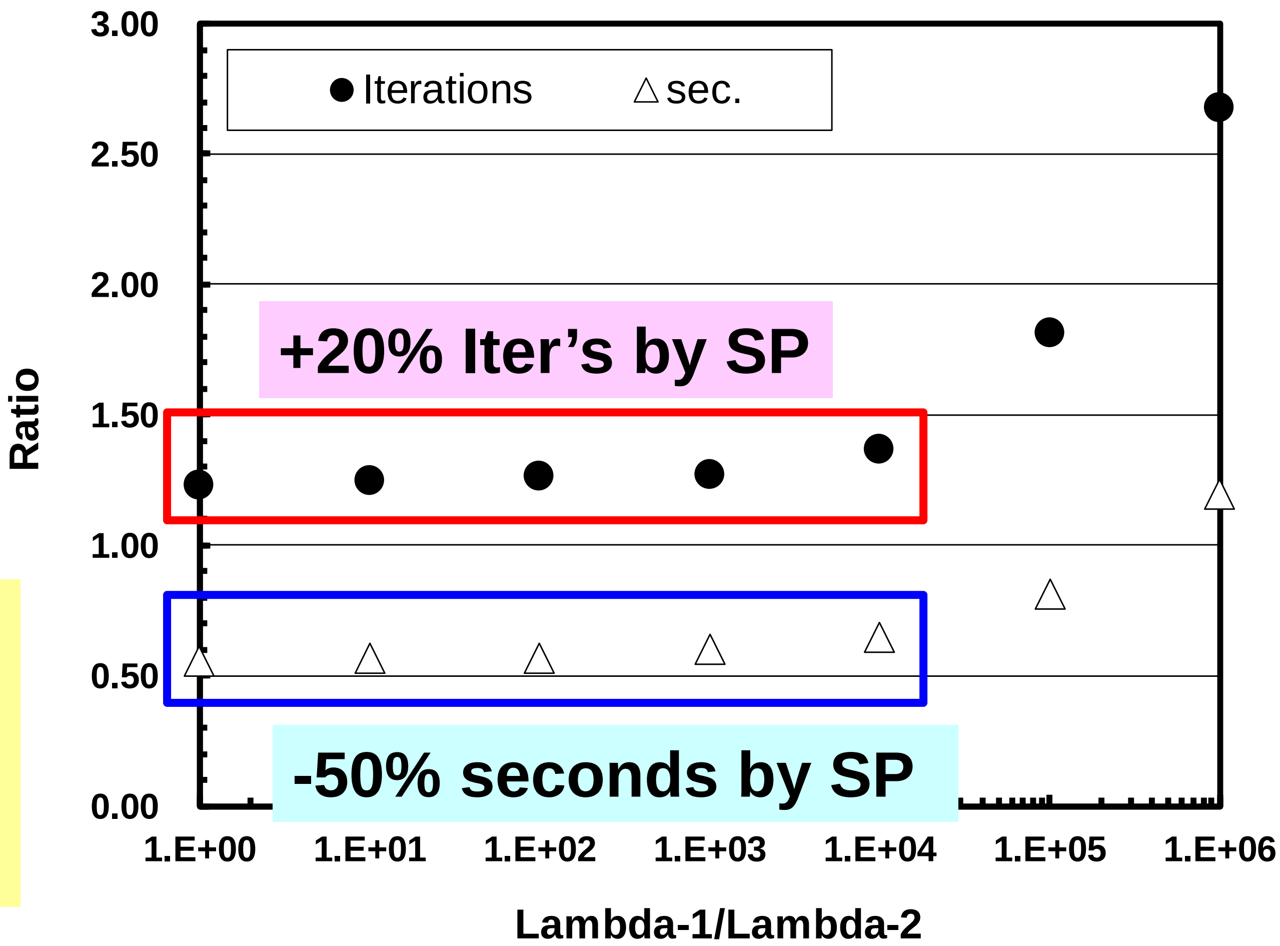

ICCG: ELL/Sliced ELL/SELL-C-σ

$$\lambda_1 = \lambda_2$$

ICCG Solvers on
Intel Xeon/Phi (KNL)
(Oakforest-PACS)
Single Node:
64/68 cores

SELL-C-σ for ICCG

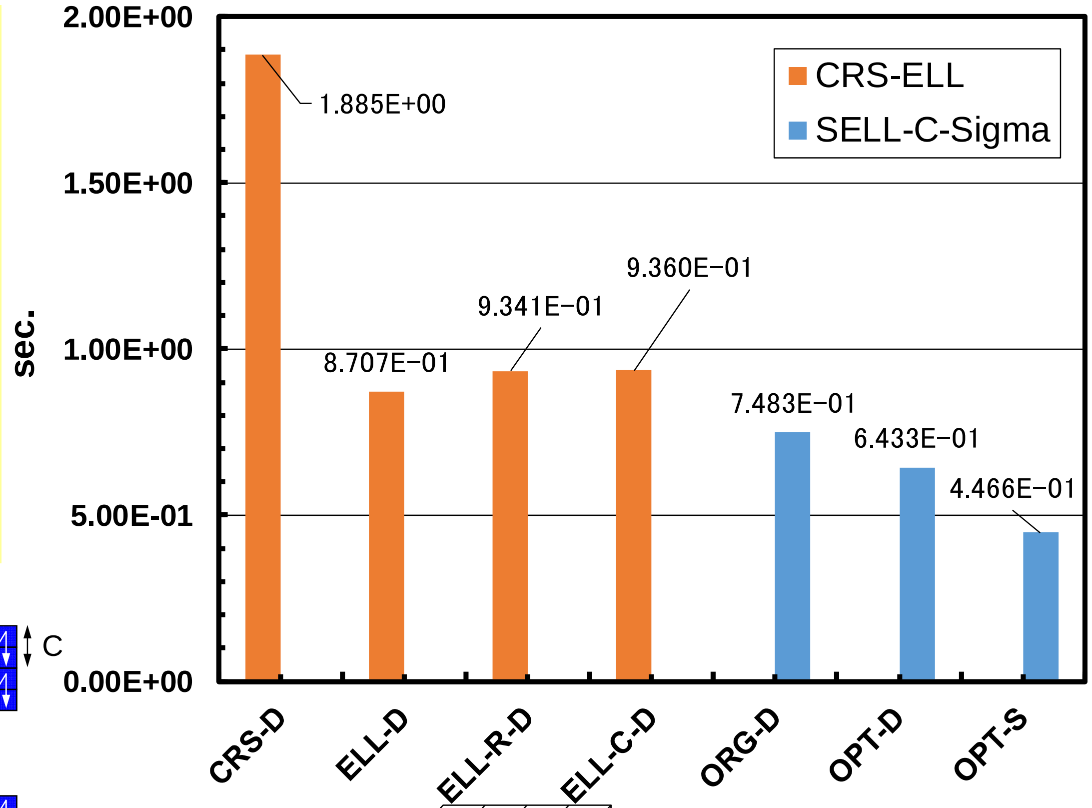

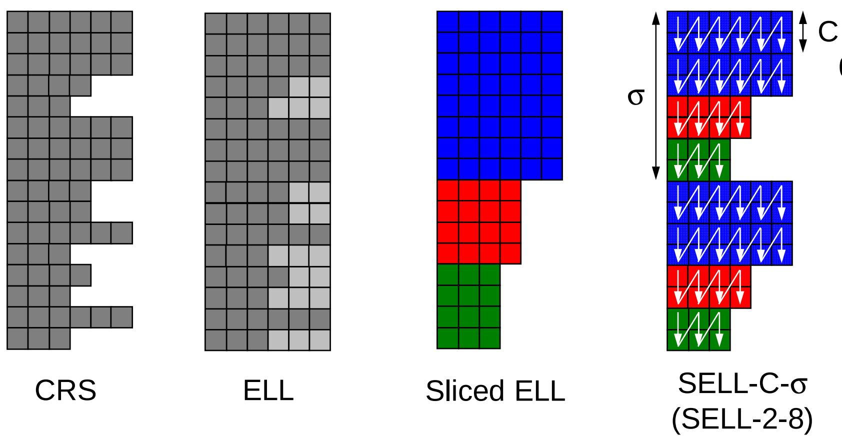

[Kreutzer, Wellein et al. 2014]

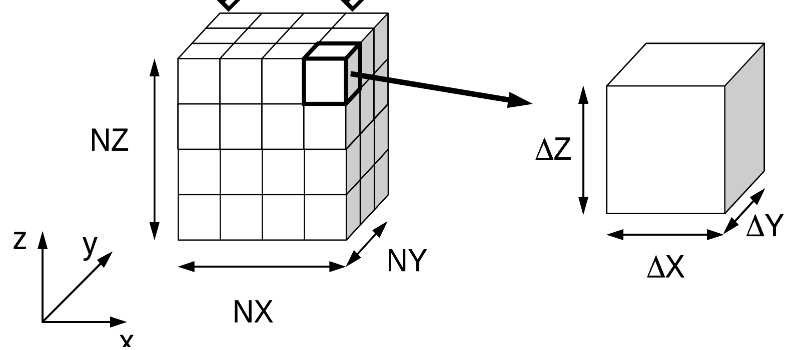

Results on OFP, Poisson-3D-OMP
Effect of SIMD Vector Length in SELL-C-σ
10 colors, 128^3

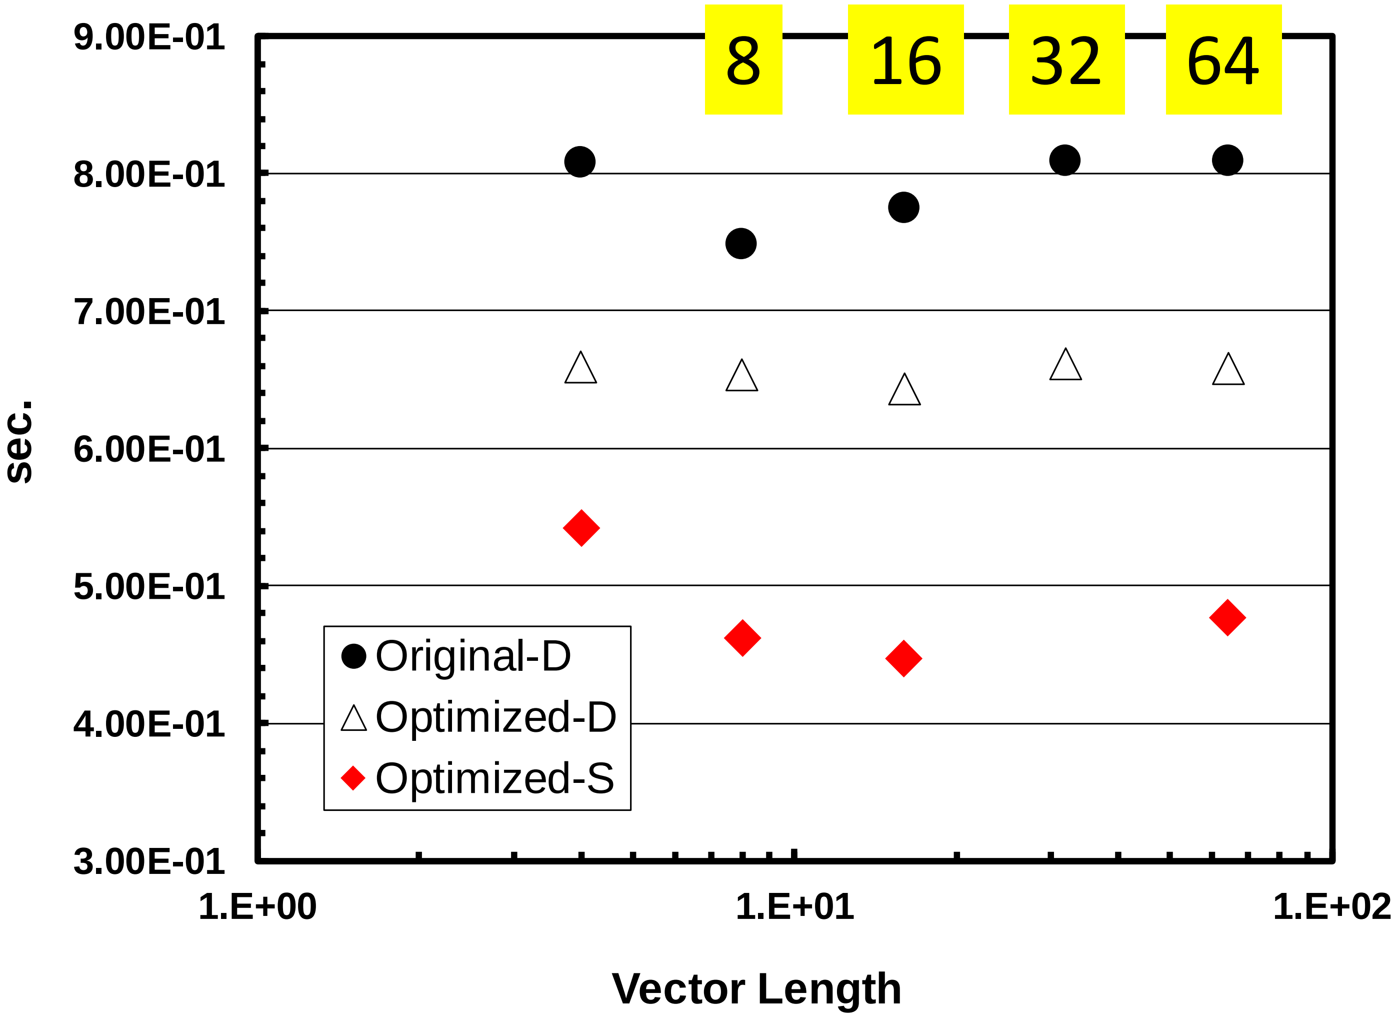

3D Poisson Solvers on Reedbush-H

$\lambda_1 = \lambda_2$

CPU only: Intel BDW: sec. & Joule

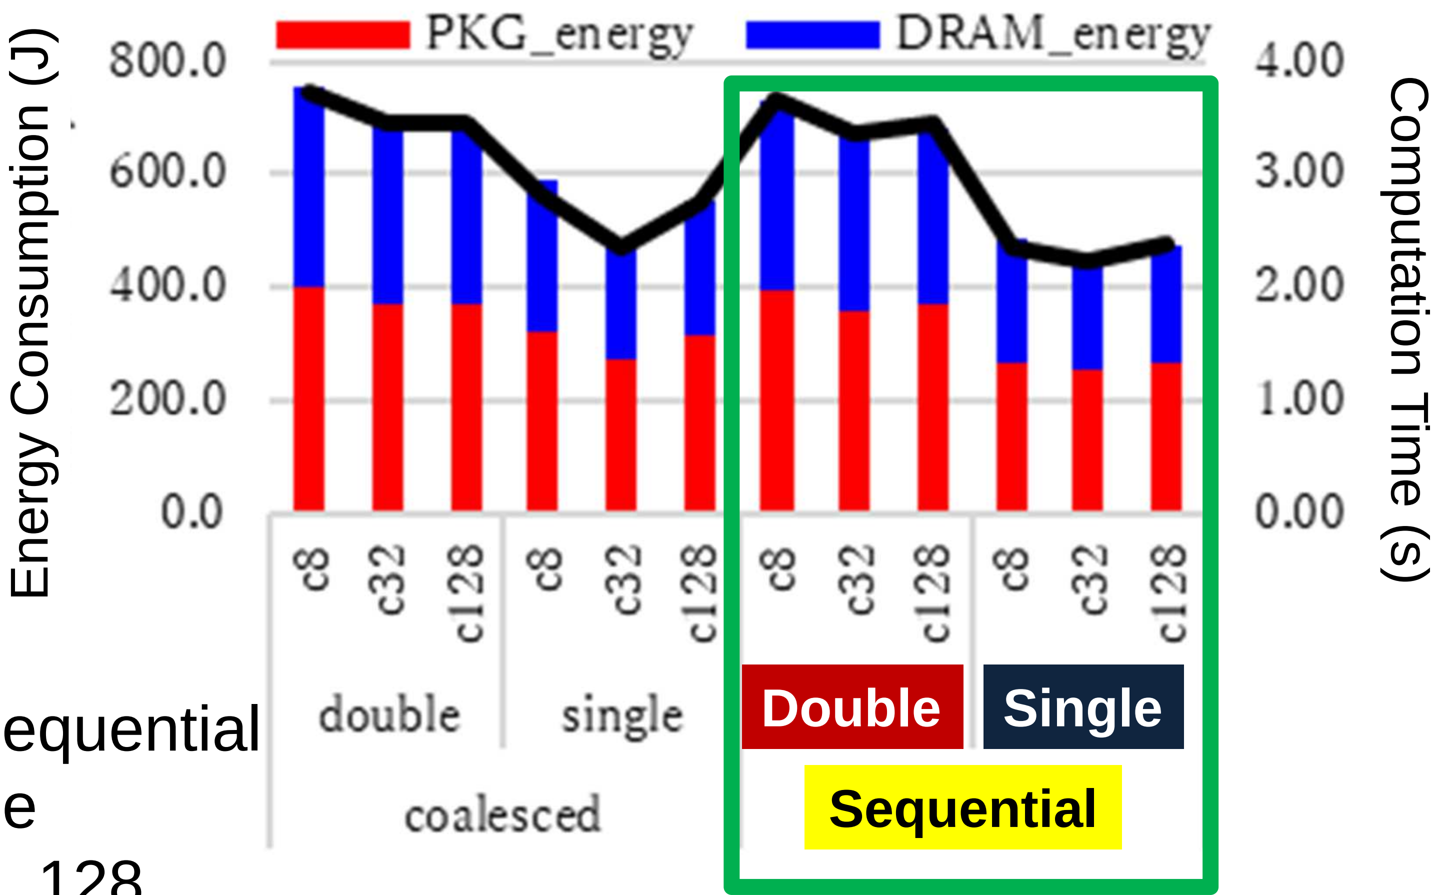

- 128^3 DOF
- Coalesced/Sequential
- Double/Single
- Colors: 8, 32, 128
- **Watt-value of SP may increase due to larger density of comp.**

[Sakamoto et al. 2018]

Computation Time (Normalized): P100, V100

[Sakamoto et al. 2018]

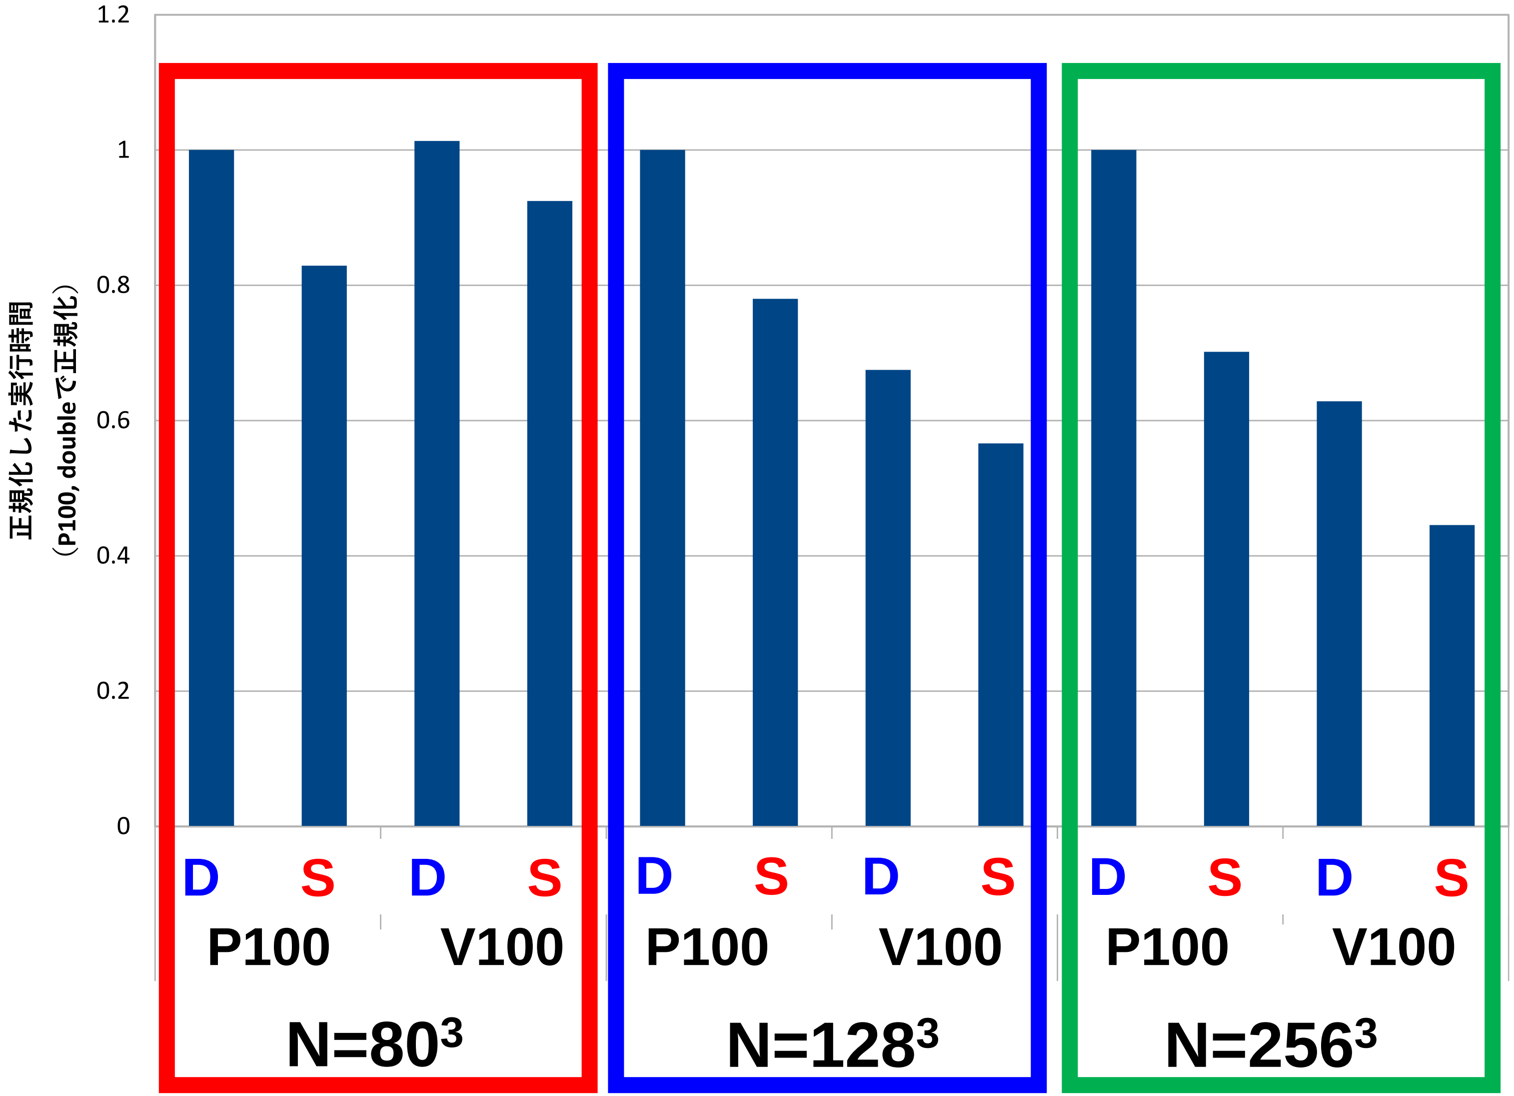

[Sakamoto et al. 2018]

Energy Consumption (Normalized): P100, V100

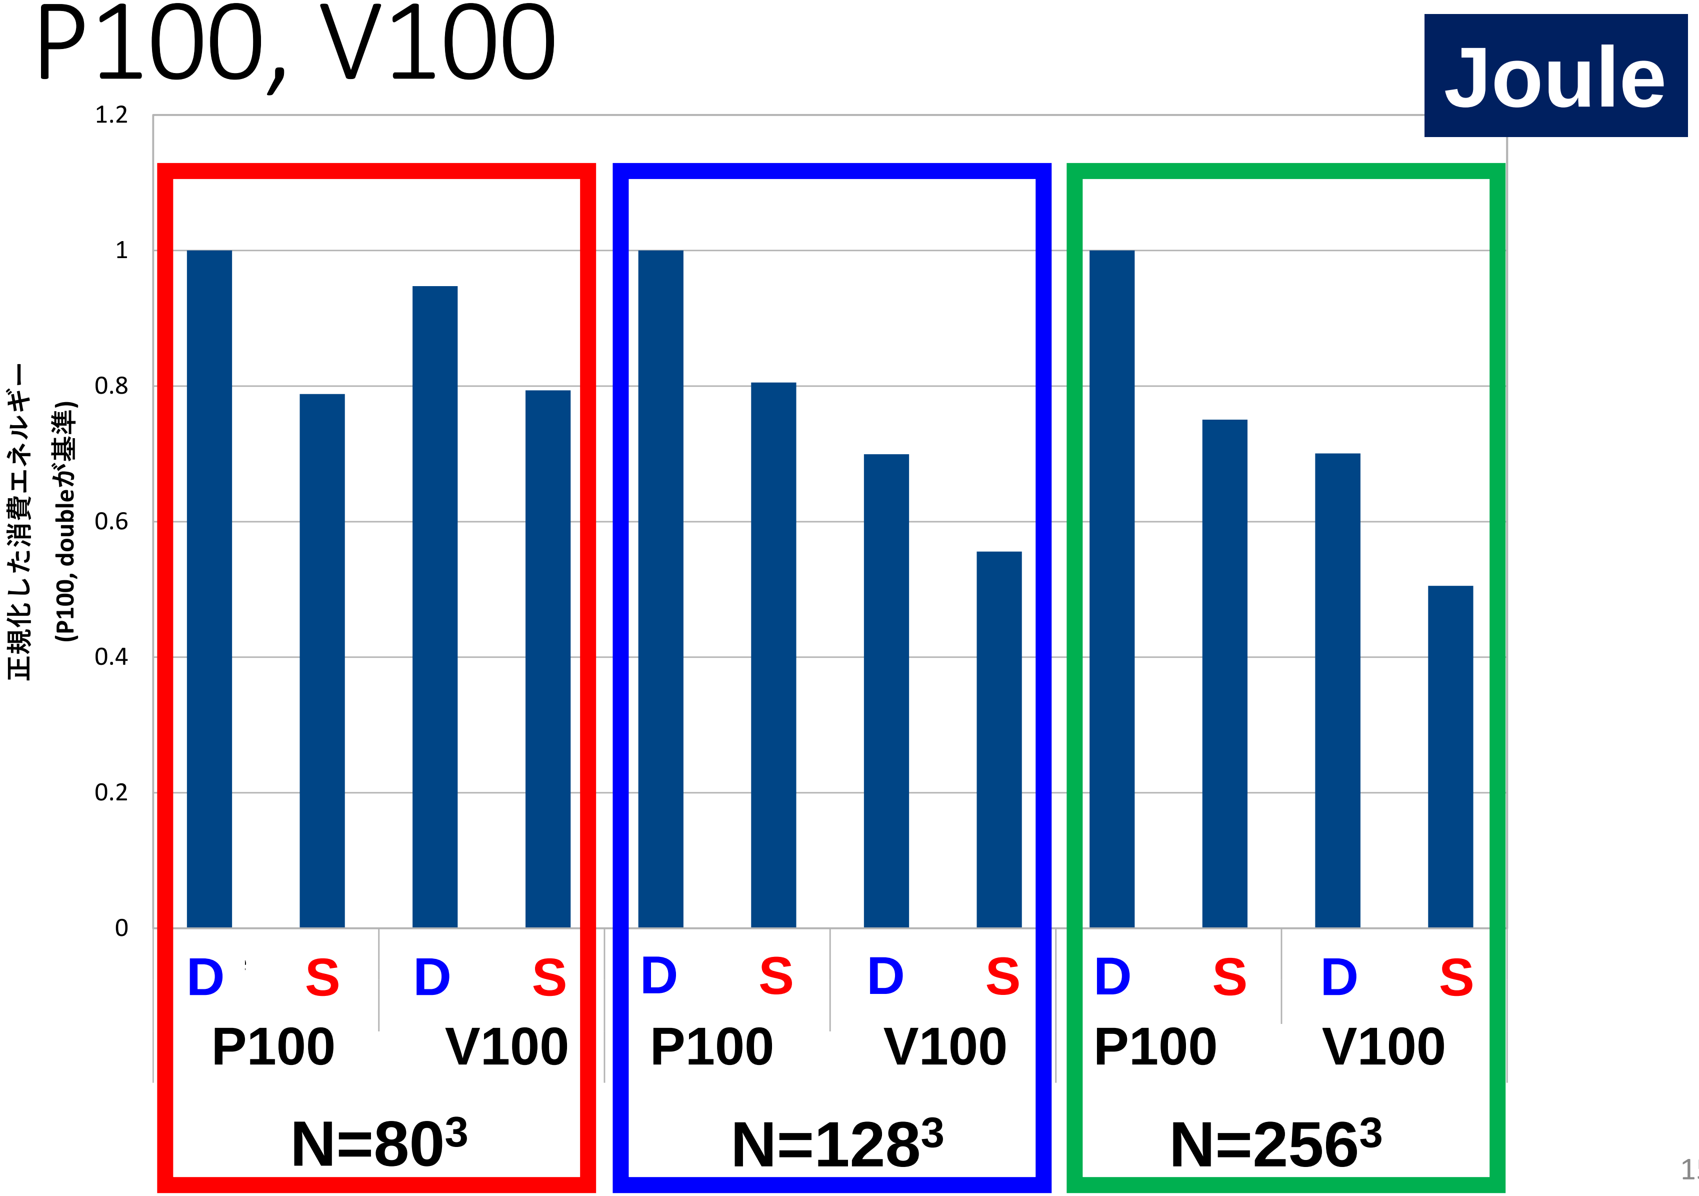

Approximate Computing with Low/Trans Precision

- Accuracy verification is important
  - Iterative Refinement
- 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)

Special Method for Rather Well-Conditioned Matrices (M-Matrix)
[Ogita, Ushiro, Oishi 2001]

Verification Algorithm

1. Solve a discretized linear system $Ax = b$.
   - $\hat{x}$: a computed solution
2. Solve a linear system $Ay = e$ where all elements of e are 1's.
   - $\hat{y}$: a computed solution
3. Verify M-property of A using $\hat{y}$. $(\hat{y} > 0 \;\Rightarrow\; A\hat{y} > 0)$
4. Compute an error bound using

$$\|x - \hat{x}\|_\infty \leq \frac{\|\hat{y}\|_\infty \|b - A\hat{x}\|_\infty}{1 - \|e - A\hat{y}\|_\infty}$$

if $\|e - A\hat{y}\|_\infty < 1$.

Processes for Verification are very similar to those of Solvers. We can do 2 processes in parallel manner simultaneously

Numerical Results

- Computer: Reedbush-U (1 node)
  - Intel Xeon E5-2695v4 (Broadwell-EP, 2.1GHz 18 cores) x 2 sockets
  - 1.21 TFLOP/s per socket, 256 GiB (153.6GB/s)
- Solver: ICCG with CM-RCM, MC(20)
- Stopping criteria:
  - For $Ax = b$, $\frac{\|b - A\hat{x}\|_2}{\|b\|_2} < 10^{-12}$
  - For $Ay = e$, $\|e - A\hat{y}\|_\infty < 10^{-2}$
- FP64 (double precision), OpenMP (36 threads)

Result (1): $\lambda_1 = \lambda_2 = 1.0$
NX=NY=NZ=128 (n = 2,097,152)

- Upper bounds of maximum relative error and relative residual norm:
  - $\max_{1 \le i \le n} \left| \frac{x_i - \hat{x}_i}{x_i} \right| \le 3.38 \times 10^{-8}$
  - $\frac{\|b - A\hat{x}\|_2}{\|b\|_2} < 3.66 \times 10^{-11}$
- Computing time

	Approximation Solve Ax=b (415 iter's)	Verification-1 Solve Ay=e (211 iter's)	Verification-2	Total
Method-1	2.38	1.18		3.56
Method-2 (2 RHS's)	2.99		1.17e-02	3.00

Result (2):

Vary $\lambda_1/\lambda_2 \sim \mathrm{cond}$ between 1 and 10^6

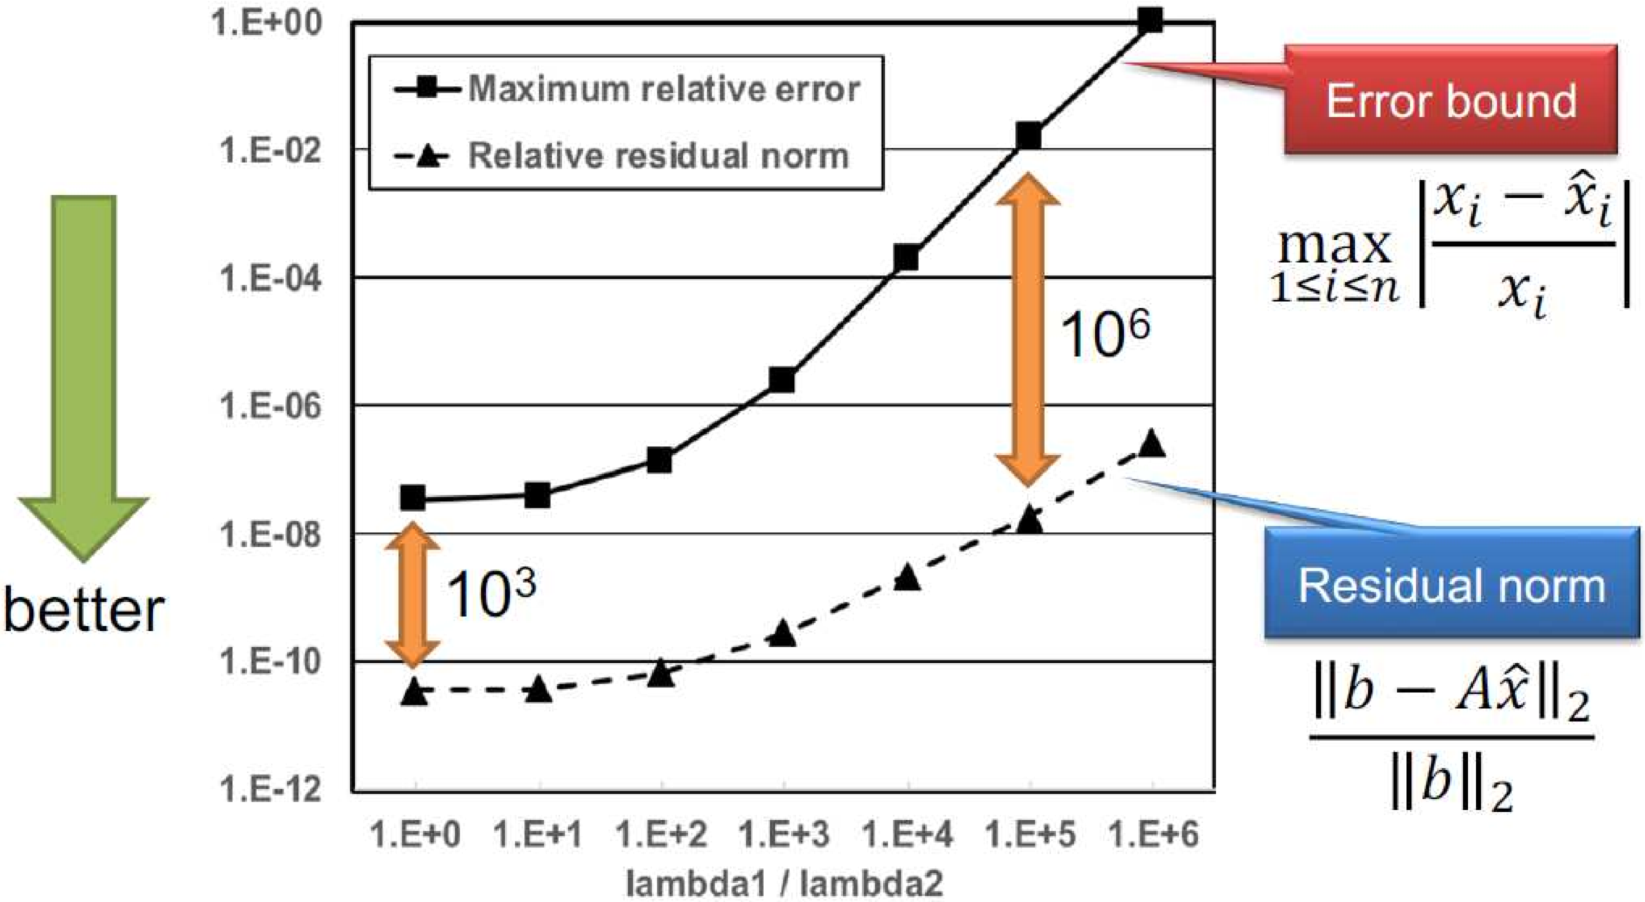

$$\max_{1 \le i \le n} \left| \frac{x_i - \hat{x}_i}{x_i} \right|$$

$$\frac{\|b - A\hat{x}\|_2}{\|b\|_2}$$

It is difficult to estimate the error of a computed solution only from residual norm!

Summary

- Numerical Library in the Exascale/Post Moore Era
  - Reduction of Energy Consumption
    - **Lower/Adaptive/Trans Precision**
    - Reduction of Computations: Data Driven Approach (DDA): Panel
- Preliminary Studies in Computing with Lower/Adaptive Precision
  - Computations with lower-precision (FP32, single precision) work for sparse matrices with certain condition number
    - Lower Power Consumption
  - Accuracy Verification
- Other Works in FY.2018
  - H-matrix solver with lower/mixed precision
  - **Iterative Refinement [Okuda 2010]**
  - Pipelined Algorithms
  - **FP16 (Half-Precision)**
    - **Severe Limitation: Only 3-digit accuracy assured**
    - **Preconditioner using Local Information: Block LU, GS**

CG Solver by Iterative Refinement

[Okuda et al. 2010]

Single*/*Double

α_i: Scaling Factor

Target Linear Equation
$Ax=b$ (Double Precision)

Coef. Matrix and RHS vector in Single Precision
$A', c_0 = b$

Inner-Loop Solver by Lower Precision
Solve $A'z_i = c_i$

New RHS of Inner Loop
$c_i = r/\alpha_i$

Iterative Refinement
$x_0 = z_0$
$x_{i+1} = x_i + \alpha_i z_{i+1}$

α_i
Tol. & Iter.# for Inner Iteration

NOT Converged
$i = i+1$

Residual
$r = b - Ax$

Converged

Convergent Solution (Double Precision) x

FP32 (Single) with FP16 Precond.

V100, All Problems converge in FP32/64

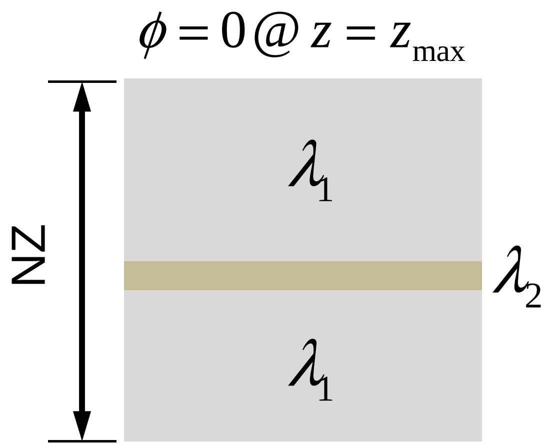

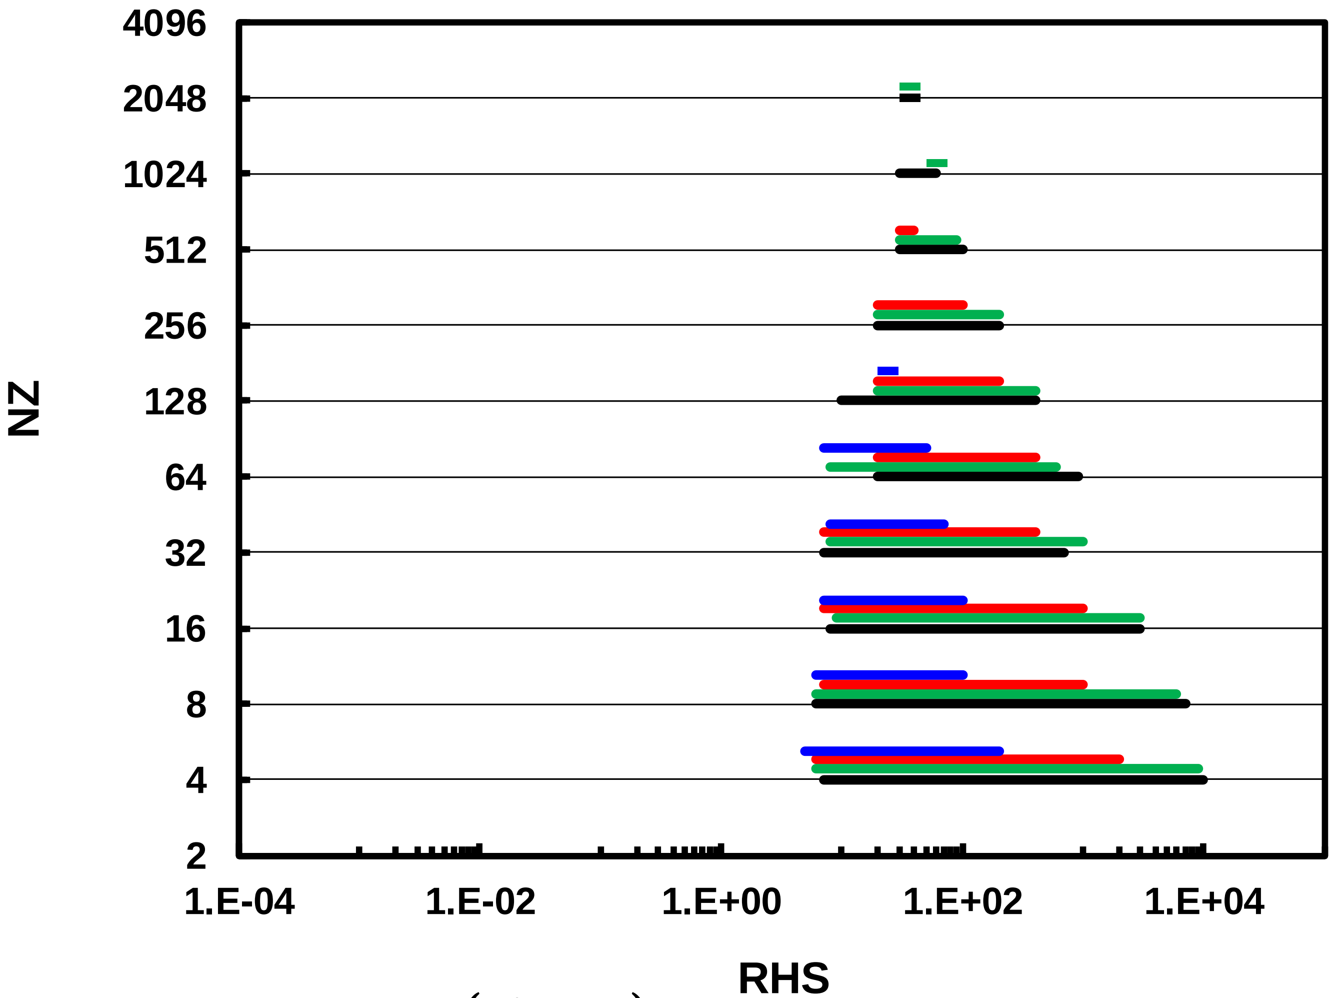

λ_1/λ_2=1.0e+3
λ_1/λ_2=1.0e+2
λ_1/λ_2=1.0e+1
λ_1/λ_2=1.0e+0

$$\nabla \cdot \left(\lambda \nabla \phi \right) = -RHS$$

[Hoshino 2018]

(Near) Future Works in FY.2019

- Accuracy Verification + AT
  - More Reasonable Method for Accuracy Verification
    - Ill-Conditioned Sparse/H Mat.: Combined with Iterative Refinement
  - Strategy for Selection of Optimum Precision
    - Accuracy, Computation Time, Power Consumption
  - Trans-Precision (e.g. FP20, FP21)
    - Challenging Approach: e.g. AT + FPGA

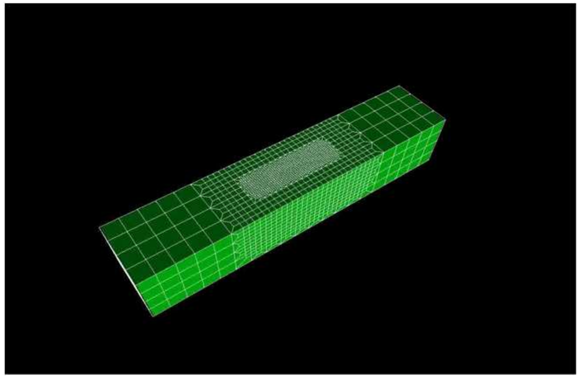

- FEM with Local Adaptive Precision
  - Precision changes on each element
    - New Idea
  - Heterogenuity, Stress Concentration, Elastic-Plastic (Linear-NL), Separation
  - Load In-Balancing in Parallel Computing
  - Discussions in WCCM 2018 in NYC

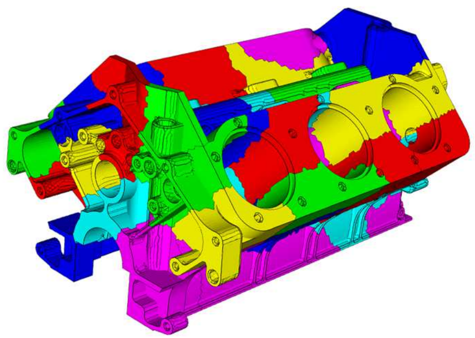

- **Towards "Appropriate Computing"**
  - **Approximate Computing + Accuracy Verification + Automatic Tuning (AT)**

Current Status

- Proposal for FY. 2019 Accepted
- Osni Marques (LBNL, USA) will join in April 2019
  – Japan-Germany-USA Collaboration
  – We welcome French collaborators !
- If you are a member, you can use:
  – Oakforest-PACS (KNL) (U.Tokyo, Tsukuba)
  – Tsubame 3 (Intel/BDW + NVIDIA P100) (Tokyo Tech)
  – Oakbridge-CX (Intel/CLX Cluster) (U. Tokyo) (After October 2019)

ICPP 2019 in Kyoto

48th International Conference on Parallel Processing

August 5-8, 2019

http://www.icpp-conf.org/

Submission Open:	**February 01, 2019**
Deadline for Submission (10-pages):	**April 15, 2019**
Author Notification:	**May 17, 2019**
Camera-Ready Due:	**June 07, 2019**




INTERNATIONAL CONFERENCE ON PARALLEL PROCESSING

Invited Speakers

Depei Qian (Sun Yat-Sen University & Beihang University, China)
Satoshi Sekiguchi (AIST, Japan)
Richard Vuduc (Georgia Tech, USA)

Please take your vacation in Japan this Summer