

# Development of a Dense Eigenvalue solver for Exascale Systems

Toshiyuki IMAMURA

RIKEN Center for Computational Science

MolSSI Workshop / ELSI Conference, the Graduate,  
Richmond VA, USA, 15-17 August 2018

# Agenda

## 1. Quick Overview of Project

- Past and Present of EigenExa
- Diagonalization of a 1million x 1million matrix on K computer

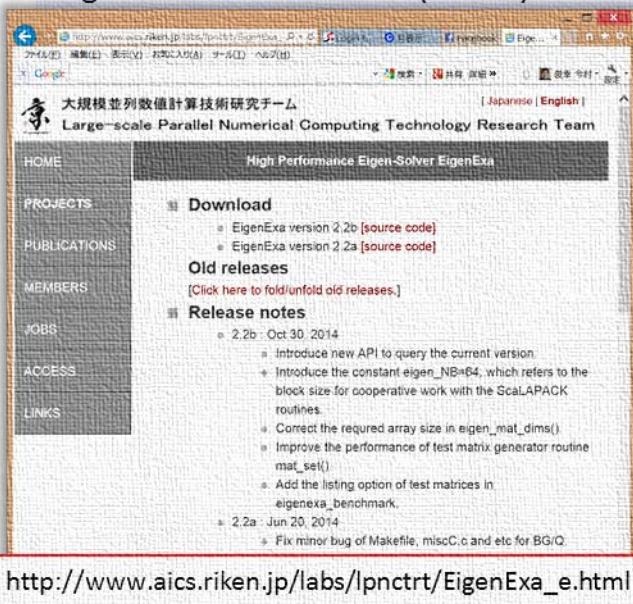
## 2. The latest updates

## 3. Future direction

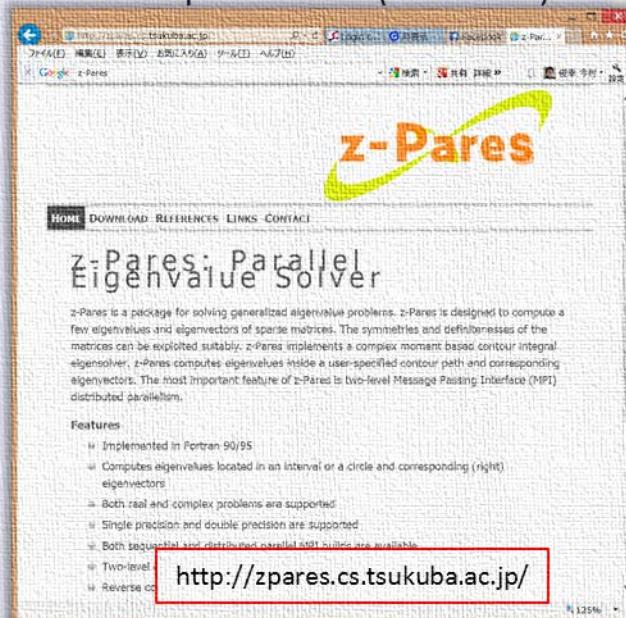
## 4. Summary

Prof. Sakurai Team 'H4ES'

EigenExa:: dense solver (RIKEN)



Z-Pares:: sparse solver (U.Tsukuba)



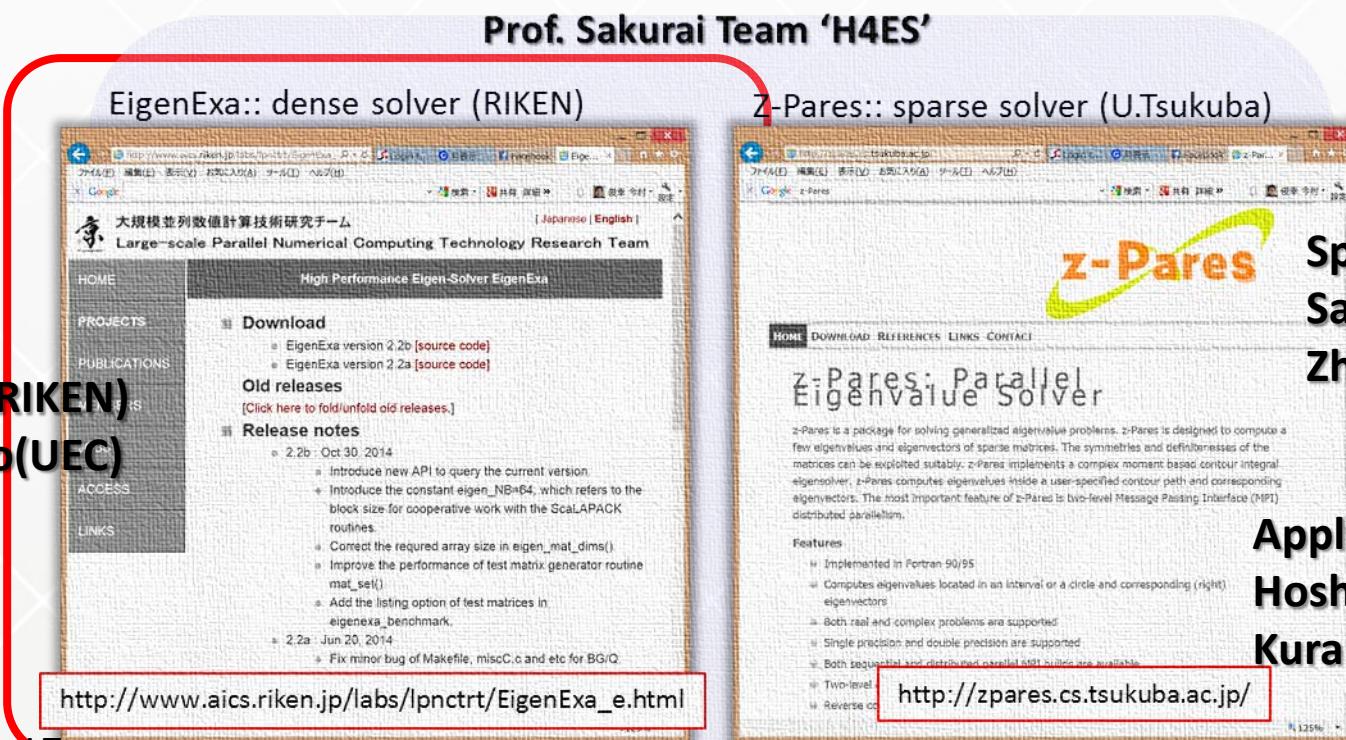
# H4ES (2011-2016, granted by CREST JST)

1. Eigenvalue problem is of significance in many scientific and engineering fields

$$Ax = \lambda Bx$$

1. Sparse and dense solver must be tightly cooperating!
2. Not only theory but computer science and applications!

**Prof. Sakurai Team 'H4ES'**



**EigenExa:: dense solver (RIKEN)**

**z-Pares:: sparse solver (U.Tsukuba)**

**Dense:**  
**Imamura(RIKEN)**  
**Yamamoto(UEC)**

**Sparse & LS:**  
**Sakurai(Tsukuba)**  
**Zhang(Nagoya)**

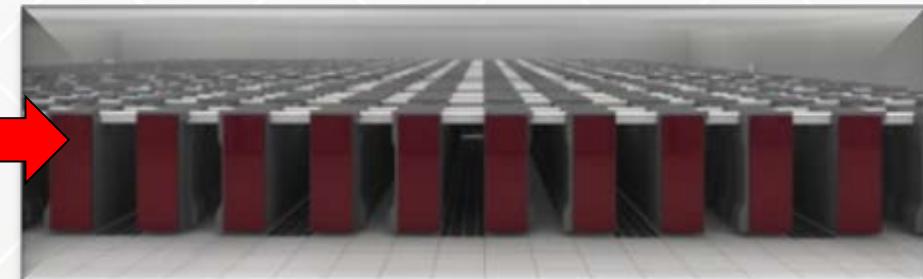
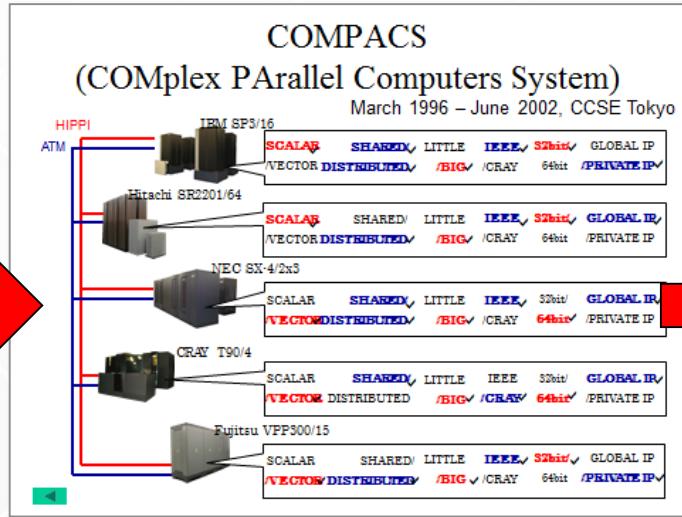
**Application:**  
**Hoshi(Tottori)**  
**Kuramashi(Tsukuba)**

[http://www.aics.riken.jp/labs/lpnctr/EigenExa\\_e.html](http://www.aics.riken.jp/labs/lpnctr/EigenExa_e.html)

<http://zpares.cs.tsukuba.ac.jp/>

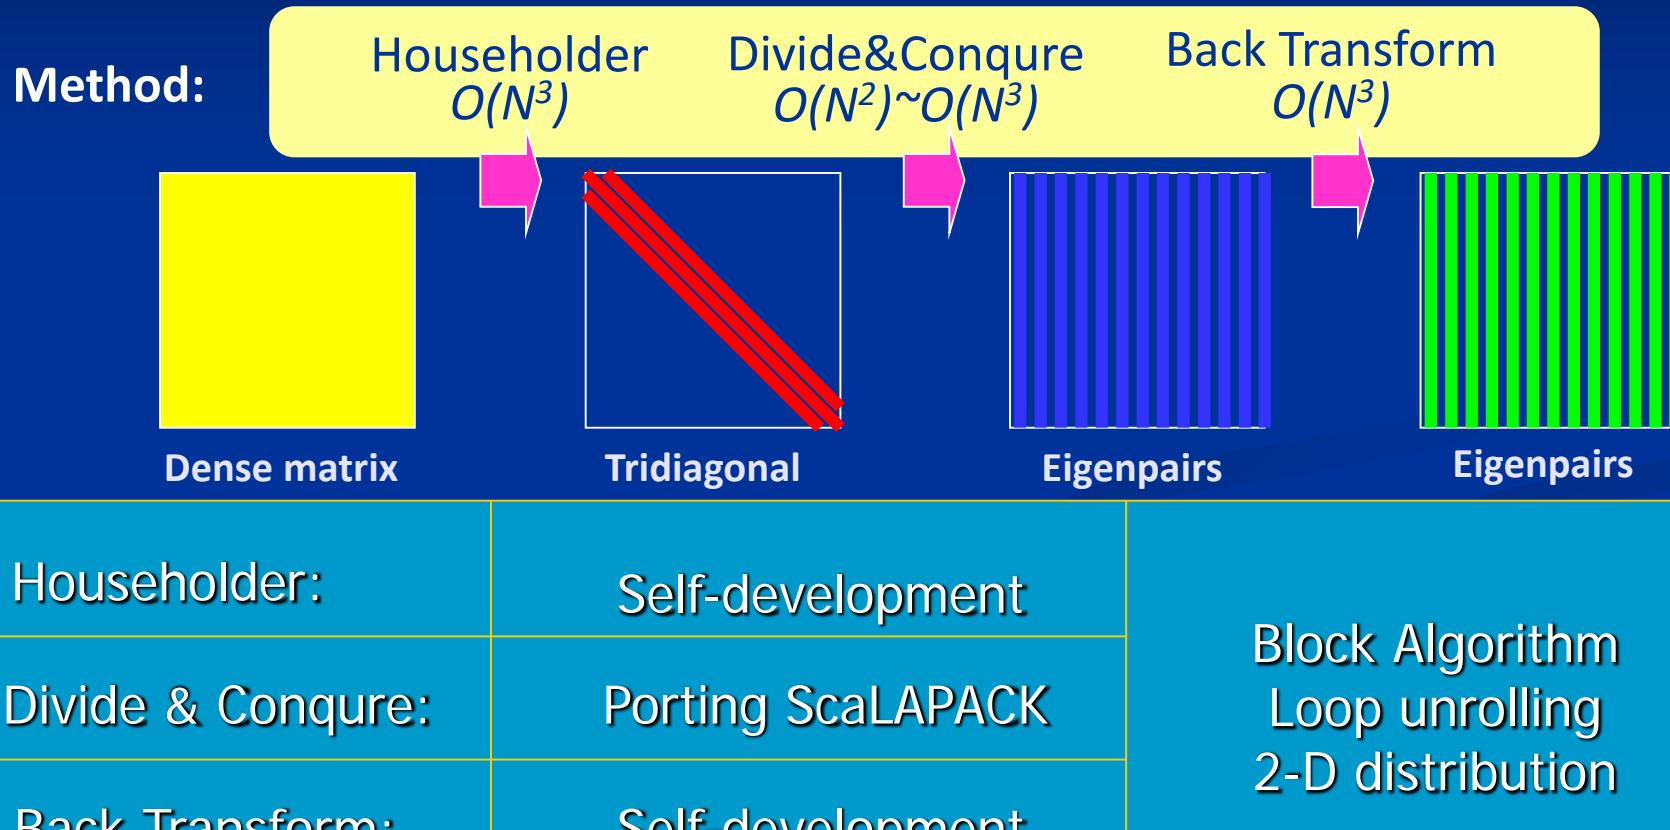
# History of the dense solver is long

## ● From VPP to the K computer



# 3, "Full Diagonalization" on ES

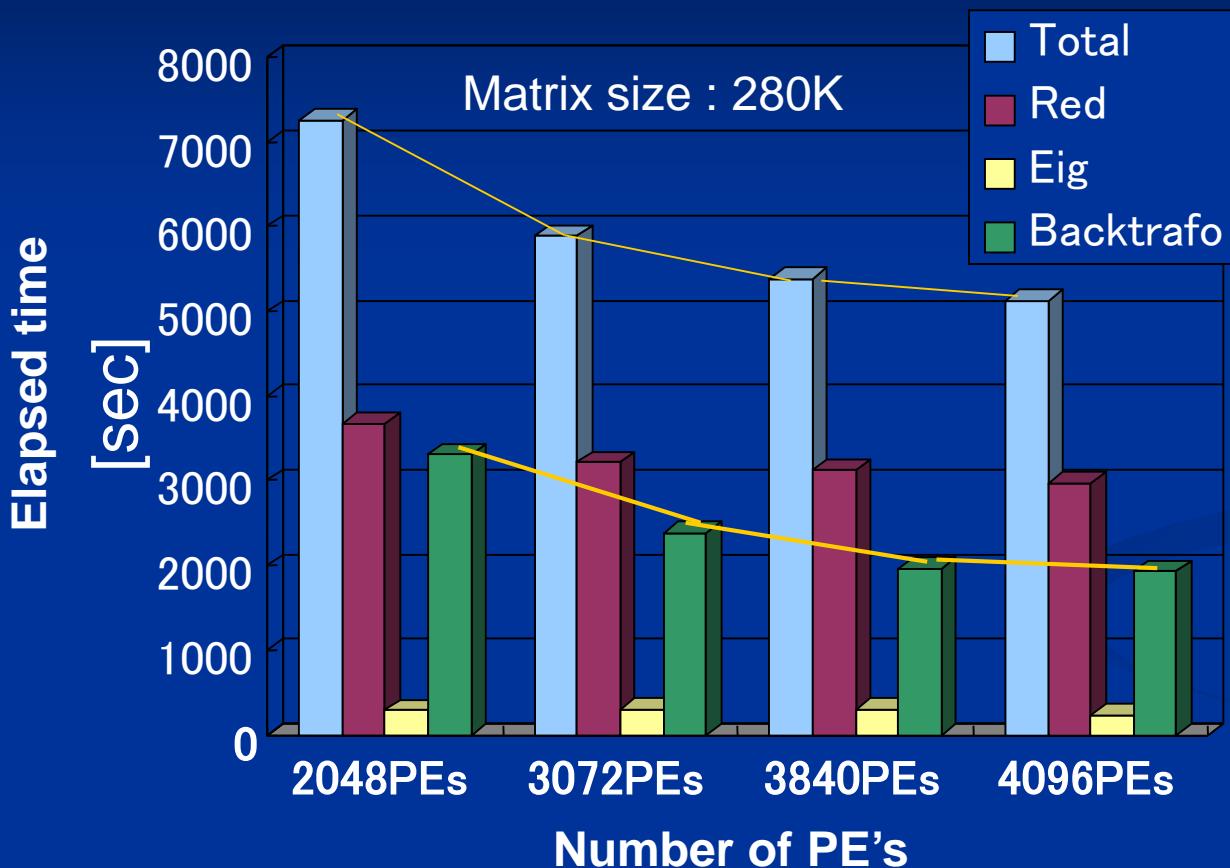
S Yamada, T Imamura, T Kano, M Machida, High-performance computing for exact numerical approaches to quantum many-body problems on the earth simulator, SC06



The best combination in terms of stability, vectorization, and parallelization

# 3, "Full Diagonalization" on ES

Test of PE No. scalability:



Matrix dim.: 280,000  
Memory: 2 T Bytes  
No. of PE's:  
2048 ~ 4096

Reduction:  
Householder transformation

Eig:  
Divide & Conquer

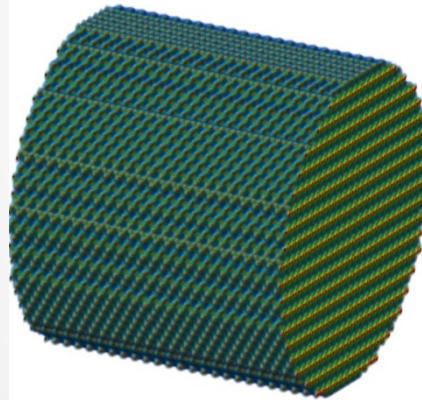
Backtrafo:  
Back transformation

**Total scalability is OK !**  
**Especially, scalability of the back transformation is good !**

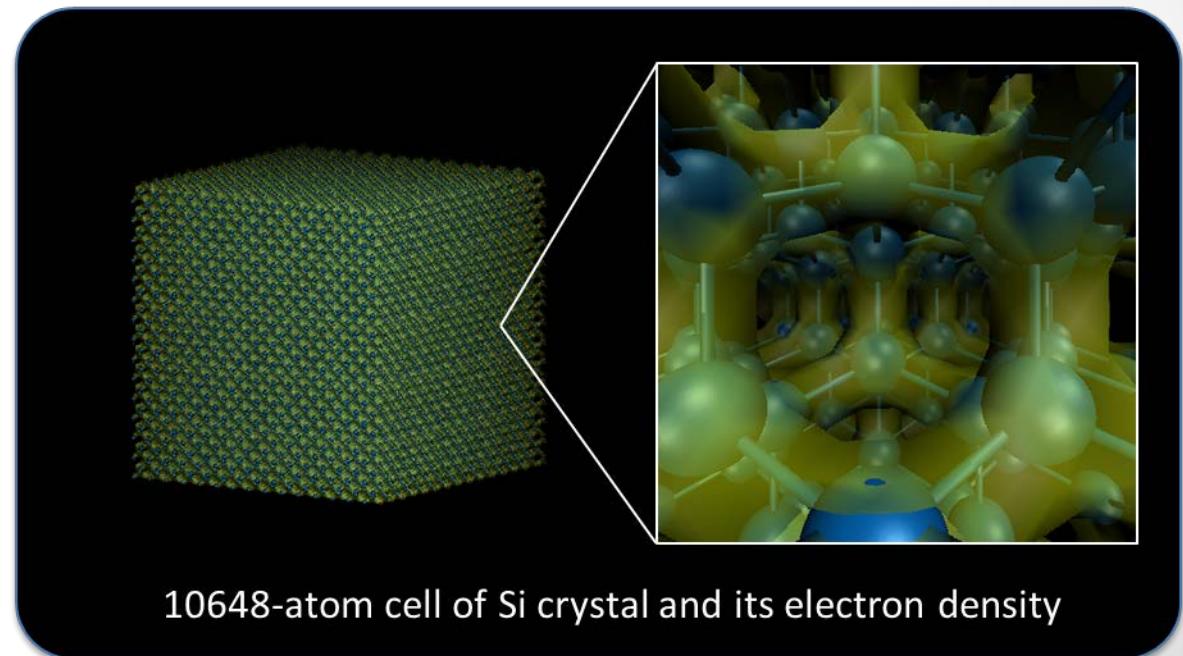


## 2.1 Review on peta-scale eigenvalue solvers

- For grand challenge applications like a nano-material simulation, they demand a high-performance eigensolver strongly.



Silicon nanowire



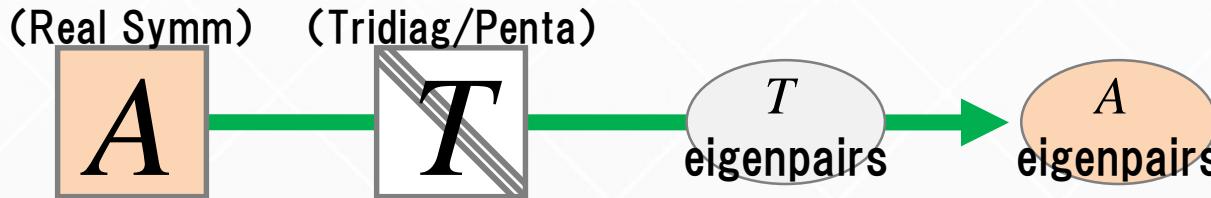
10648-atom cell of Si crystal and its electron density

*Y.Hasegawa, et al @AICS.RIKEN, GBP winner, SC2011*

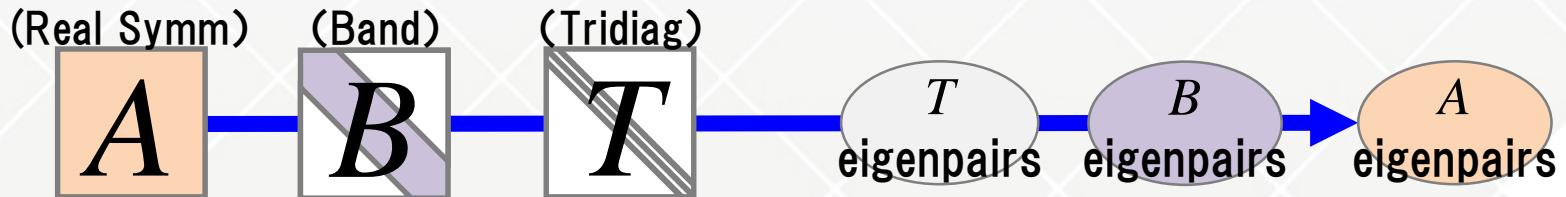


# Two major algorithms for dense SEVP

- **1-stage algorithm (EigenExa, ELPA1, Wilkinson512/Kevd)**
  - Standard numerical scheme
  - Householder-tridiagonalization has memory-bound bottleneck



- **2-stage algorithm (ELPA2, PLASMA, MAGMA)**
  - Transform to Banded-matrix = Less Byte/Flop operation by GEMM
  - Needed 2-stage back-transformation, but slows unfortunately. However, a big advantage when a few eigenmodes are computed



# Divide and Conquer method for a banded matrix

$(\Lambda, X) = \text{Divide and Conquer}(\text{matrix } B)$

Decompose B into  $B_1$  and  $B_2$  with k-perturbations as

$$B = B_1 \oplus B_2 + \alpha_1 c_1 c_1^t + \alpha_2 c_2 c_2^t + \dots + \alpha_b c_b c_b^t$$

$(\Lambda_1, X_1) = \text{Divide and Conquer } (B_1);$   
 $(\Lambda_2, X_2) = \text{Divide and Conquer } (B_2);$

$$B^{(1)} = \Lambda^{(0)} + \alpha_1 d_1 d_1^t, \quad d_1 = X^{(0)^t} c_1$$

Solve eigenvalues and vectors of  $B^{(1)}$  via secular eq.

$$f(\lambda) = 1 + \alpha_1 \sum_{i=1}^n \frac{d_1^2}{\lambda_i - \lambda} = 0 \quad \rightarrow \Lambda^{(1)}$$

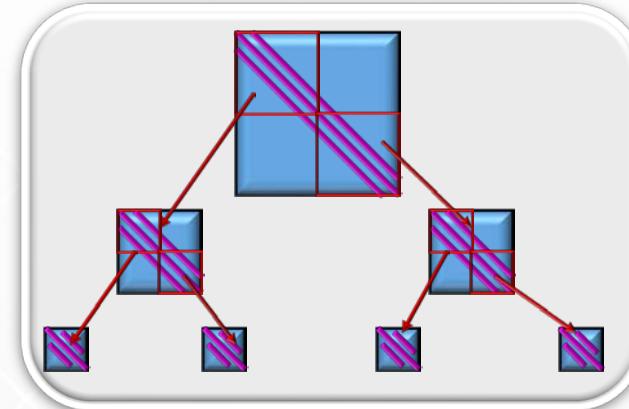
$$x = (X_1, X_2) \text{Normal}((\Lambda^{(0)} - \lambda I)^{-1} d_1) \quad \rightarrow X^{(1)}$$

Do i=2,...

$$B^{(i)} = \Lambda^{(i-1)} + \alpha_i d_i d_i^t, \quad d_i = \prod X^{(i-1)^t} c_i$$

$$\Lambda = \Lambda^{(k)} \\ X = \prod X^{(k)}$$

Rough estimation: the cost  $O((2b-1)N^3)$ ,  
it increases  $2k-1$  times compared with  $k=1$ .  
Usually, New 1-step scheme shows better  
performance than old-1-step or 2-step.

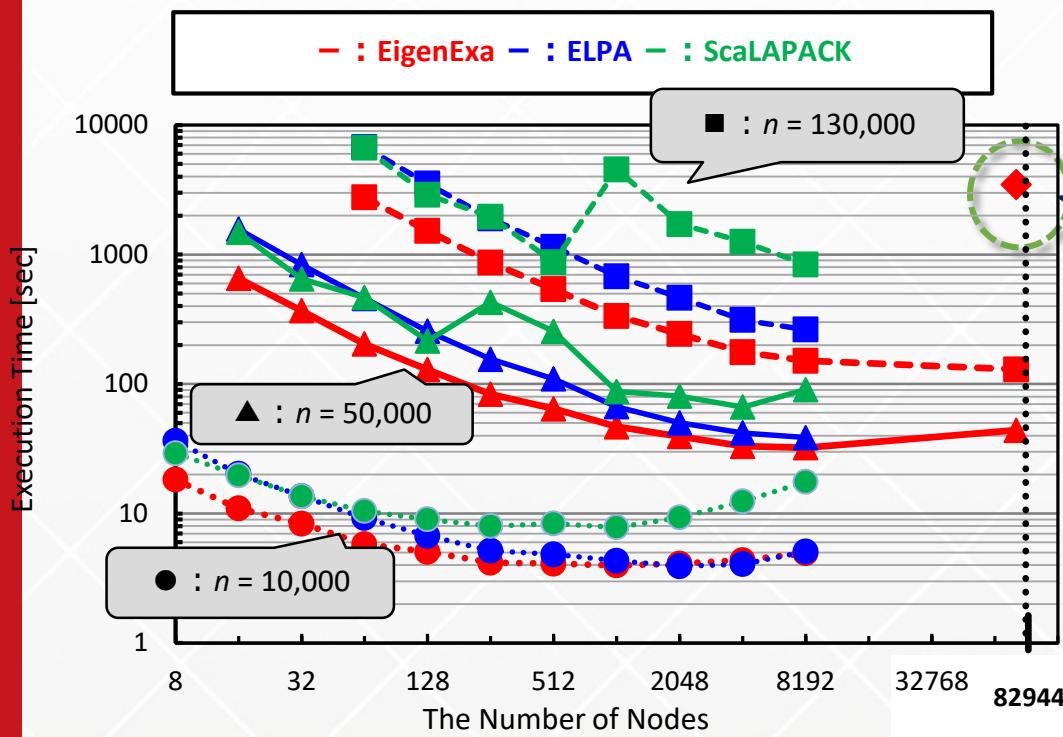


\*This type of DC algorithm was discussed by Arbenz, Gander, Golub, Gansterer, et al. in early 1990's.

\*We implemented a parallel version adopting a robust technique to modify orthogonality of eigenvectors.

# EigenExa: World Largest Dense Eigenvalue Computation

- We have successfully done a world largest-scale dense eigenvalue benchmark (one million dimension) by EigenExa taking advantage of the overall nodes (82,944 processors) of K computer in 3,464 seconds. Our EigenExa achieves 1.7 PFLOPS (16% of the K computer's peak performance).
- Feasibility and Reliability for the algorithm and the library are confirmed, especially assumed on a post-K system.**



◆ :  $n = 1,000,000$

EigenExa solves a world largest-scale problem.

(1.7 PFLOPS, 16% of K computer's theoretical peak performance)

$$\max_i \|Av_i - \lambda_i v_i\|_2 / \|A\|_F = 3.1 \times 10^{-13}$$
$$\|V^T V - I\|_F = 2.1 \times 10^{-10}$$

- ✓  $n$  is the dimension of problems.
- ✓ 1 MPI process \* 8 threads per node.
- ✓ Test matrices are randomly generated.

## Specification of K computer

- Peak performance: 10.6 PFLOPS
- Num. of Nodes: 82,944
- Performance/node: 128 GFLOPS (One octa-core SPARC 64 VIIIfx)
- Network: Tofu interconnect (6D mesh-torus)

Related performance report is  
T.Fukaya, Tl. "Performance evaluation of the EigenExa eigensolver on Oakleaf-FX:  
tridiagonalization versus pentadiagonalization", PDSEC2015

# Agenda

## 1. Quick Overview of Project

- Past and Present of EigenExa
- Diagonalization of a 1million x 1million matrix on K computer

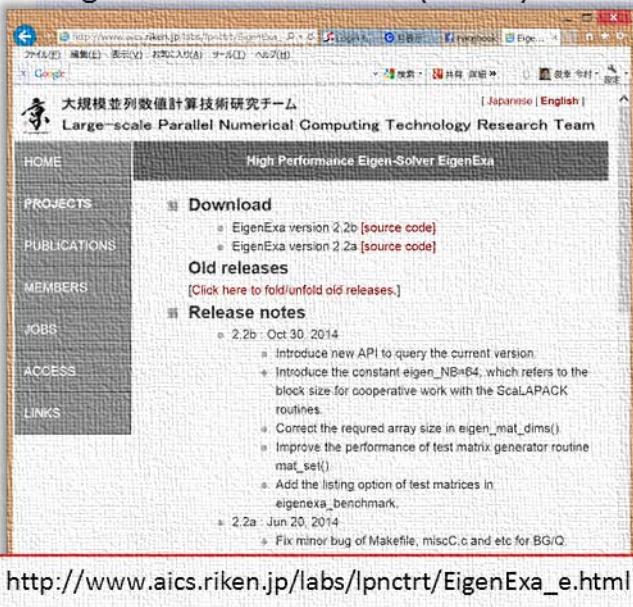
## 2. The latest updates

## 3. Future direction

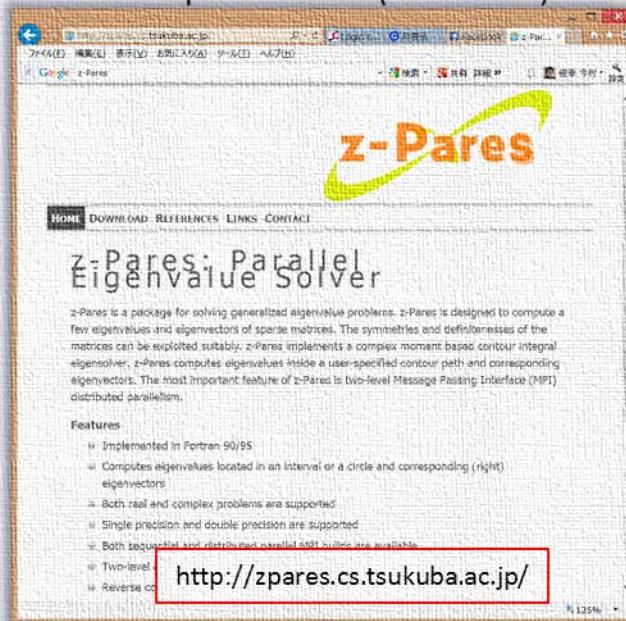
## 4. Summary

Prof. Sakurai Team 'H4ES'

EigenExa:: dense solver (RIKEN)



Z-Pares:: sparse solver (U.Tsukuba)



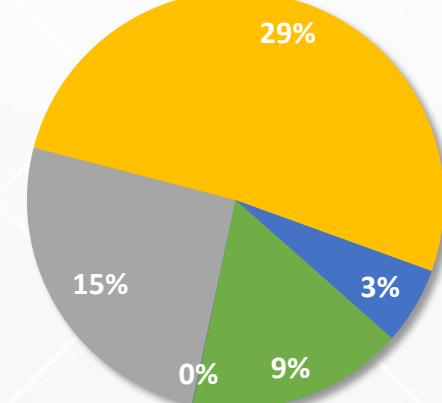
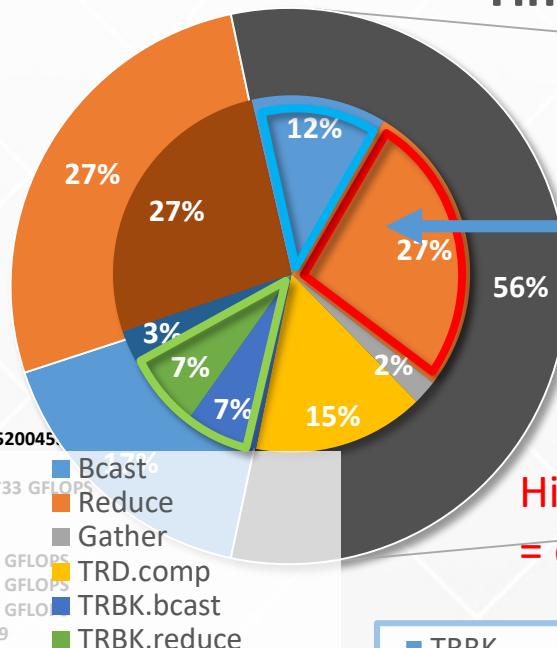
# Bottleneck is network

```

NUM.OF.PROCESS= 82944 ( 288 288 )
NUM.OF.THREADS= 8
calc (u,beta) 503.0970594882965
mat-vec (Au) 1007.285000801086 661845.1244051798
2update (A-uv-vu) 117.4089198112488
5678160.294281102
calc v 0.0000000000000000
v=v-(UV+VU)u 328.3385872840881
UV post reduction 0.6406571865081787
COMM_STAT
  BCAST :: 424.3022489547729
  REDUCE :: 928.1299135684967
  REDIST :: 0.0000000000000000
  GATHER :: 78.28400993347168
TRD-BLK 1000000 1968.435860157013
677356.7583893638 GFLOPS
TRD-BLK-INFO 1000000 48
before PDSTEDC 0.1448299884796143
PDSTEDC 905.2210271358490
MY-REDIST1 1.544256925582886
MY-REDIST2 14.75343394279480
RERE1 4.861211776733398E-02
COMM_STAT
  BCAST :: 4.860305786132812E-02
  REDUCE :: 2.155399322509766E-02
  REDIST :: 0.0000000000000000
  GATHER :: 0.0000000000000000
PDGEMM 532.6731402873993 5417097.56520045
GFLOPS
D&C 921.8044028282166 3130319.580211733 GFLOPS
TRBAK= 573.9026420116425
COMM= 533.7601048946381
  573.9026420116425 3484911.644577213 GFLOPS
  182.3303561210632 5484550.248648792 GFLOPS
  152.0370917320251 6577342.335399065 GFLOP
  0.1022961139678955 7.379654884338379
COMM_STAT
  BCAST :: 229.3666801452637
  REDUCE :: 234.4477448463440
  REDIST :: 0.0000000000000000
  GATHER :: 0.0000000000000000
TRBAKWy 573.9029450416565
TRDBAK 1000000 573.9216639995575 3484796.141101135
GFLOPS
Total 3464.162075996399 1795203.448396145 GFLOPS
Matrix dimension = 1000000
Internally required memory = 480502032 [Byte]
Elapsed time = 3464.187163788010 [sec]
  
```

## The World Largest Dense Eigenvalue Computation

### Time Breakdown



Highlighted pie  
= communication

■ TRBK	■ D&C	■ TRD.Reflector
■ TRD.AU(MVs)	■ TRD.2k-update	■ TRD.ComputeV
■ TRD.Local update		

# CA for EigenExa

- Communication Avoiding for Householder transformation unlike CAQR

- Blocking technique, increasing locality by data replication, and exchange the operation order.
- Introducing an extended form of vector 'A'.
- Computing  $Au$  and  $u^T u$ , simultaneously.

TI. "Communication-Avoiding approaches of dense Eigenvalue / SVD problems", PMAA18

TI. "Parallel dense eigenvalue solver and SVD solver for post-petascale computing systems ", PMAA16

TI, etc, "CAHTR: Communication-Avoiding Householder Tridiagonalization", ParCo15

**Principles : Distributive Law && Exchange order && Introducing the correction terms && Combine couples of collective ops.**

- naive

$$s = \text{sign}(\|u\|, -(u, e))$$

$$u := u - se$$

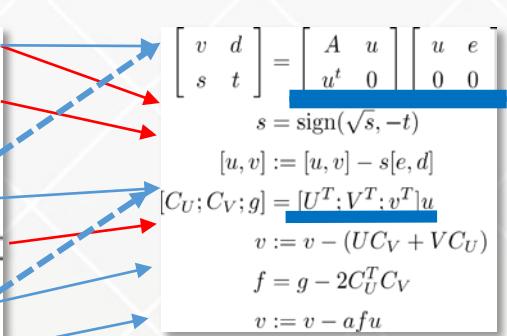
$$v = Au$$

$$[C_U; C_V] = [U^T; V^T]u$$

$$v := v - (UC_V + VC_U)$$

$$f = (u, v)$$

$$v := v - afu$$



- optimal

$$\begin{array}{c|ccccc} v & d & & & & \\ \hline s & t & & & & \\ C_U & \gamma_U & & & & \\ C_V & \gamma_V & & & & \\ \hline g & v_1 & & & & \\ v_1 & a_{11} & & & & \end{array} = \begin{array}{c|cccc} I & 0 & 0 & 0 & \\ \hline 0 & I & 0 & 0 & \\ 0 & 0 & I & 0 & \\ 0 & 0 & 0 & I & \\ \hline u^t & 0 & 0 & 0 & \\ e^t & 0 & 0 & 0 & \end{array} \begin{array}{c|cc} A & u & U & V \\ \hline u^t & 0 & 0 & 0 \\ U^t & 0 & 0 & 0 \\ V^t & 0 & 0 & 0 \end{array} \begin{array}{c|cc} u & e \\ \hline 0 & 0 \\ 0 & 0 \\ 0 & 0 \end{array}$$

$$s = \text{sign}(\sqrt{s}, -t)$$

$$[u, v] := [u, v] - s[e, d]$$

$$[C_U; C_V] = [C_U; C_V] - s[\gamma_U; \gamma_V]$$

$$v := v - (UC_V + VC_U)$$

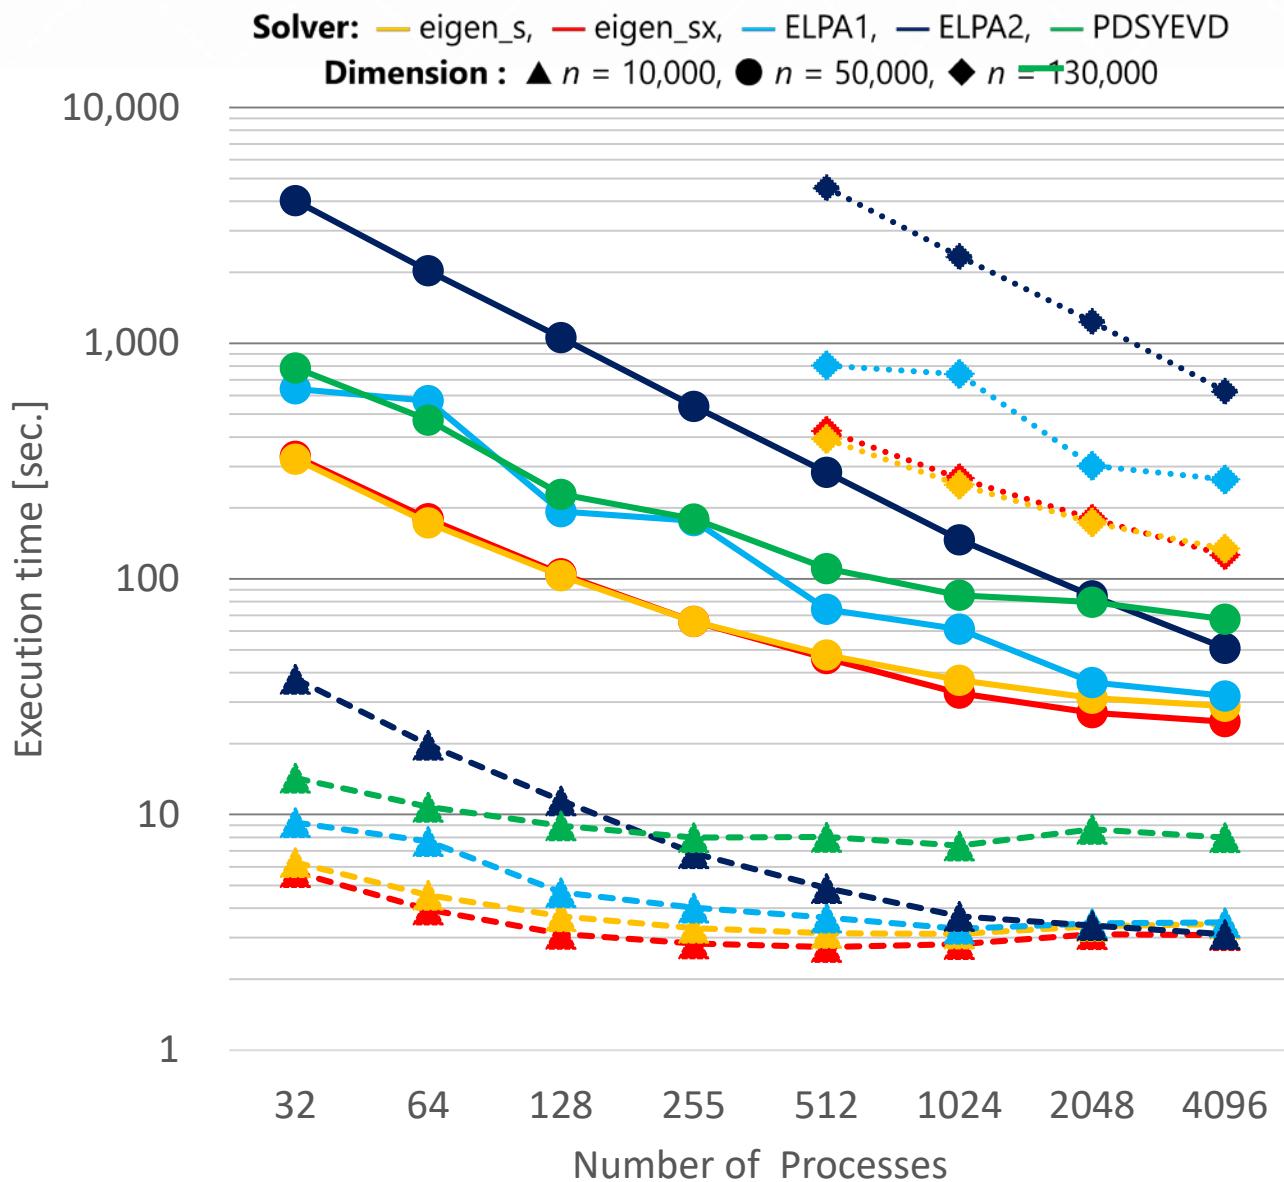
$$f = g - 2C_U^T C_V - s(2v_1 - sa_{11})$$

$$v := v - afu$$

# Performance of the latest EigenExa (2.5RC)

- Test configurations:

- K computer
- One SPARC64 VIIIfx CPU (128 GFLOPS, 8 cores) / node
- Tofu (6-dimensional torus internode network)
- all eigenpairs of  $n \times n$  Frank matrix
- 1 MPI process/node, 8 threads/process
- Fujitsu multithreaded BLAS / MPI libraries
- ELPA version 2016
- Without openmp for loop optimization but thread-parallelized BLAS is used.



# Agenda

## 1. Quick Overview of Project

- Past and Present of EigenExa
- Diagonalization of a 1million x 1million matrix on K computer

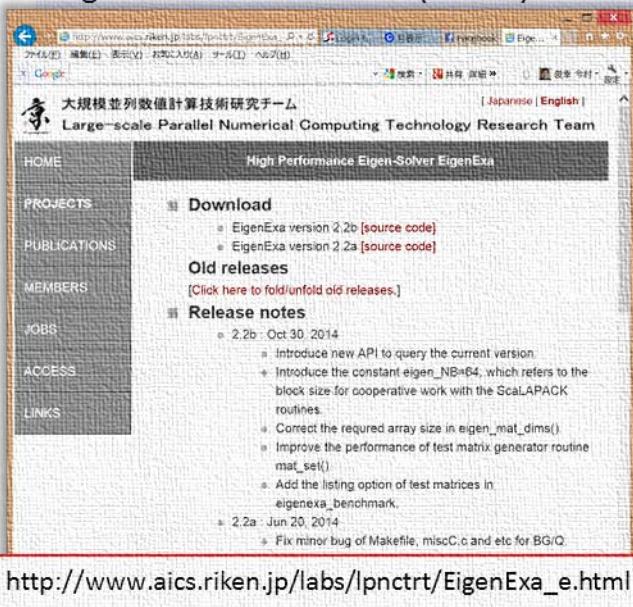
## 2. The latest updates

## 3. Future direction

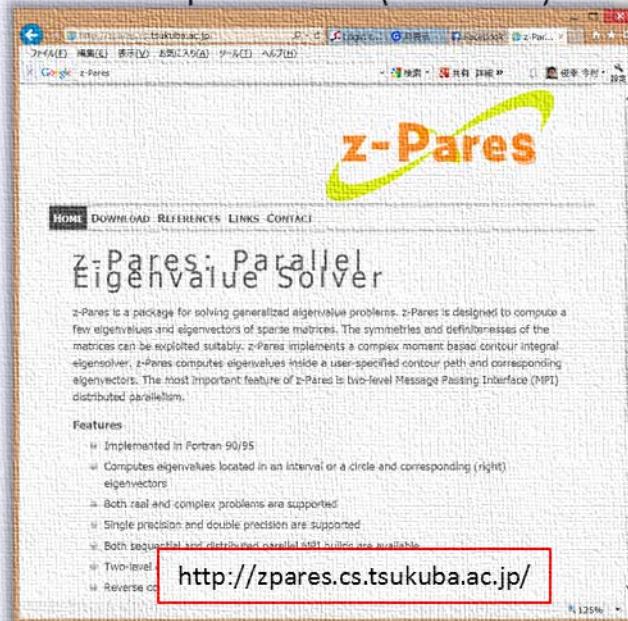
## 4. Summary

Prof. Sakurai Team 'H4ES'

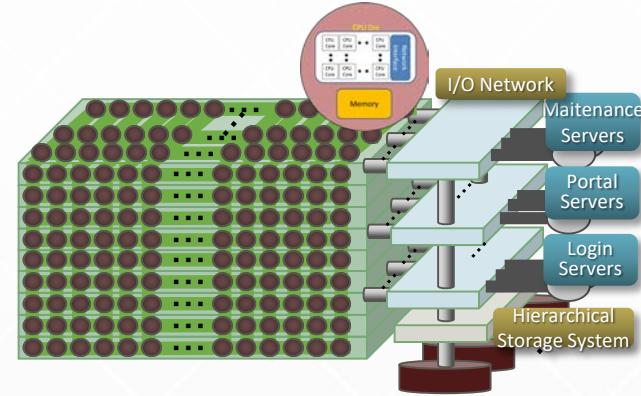
EigenExa:: dense solver (RIKEN)



Z-Pares:: sparse solver (U.Tsukuba)



# FLAGSHIP2020 Project



## □ Missions

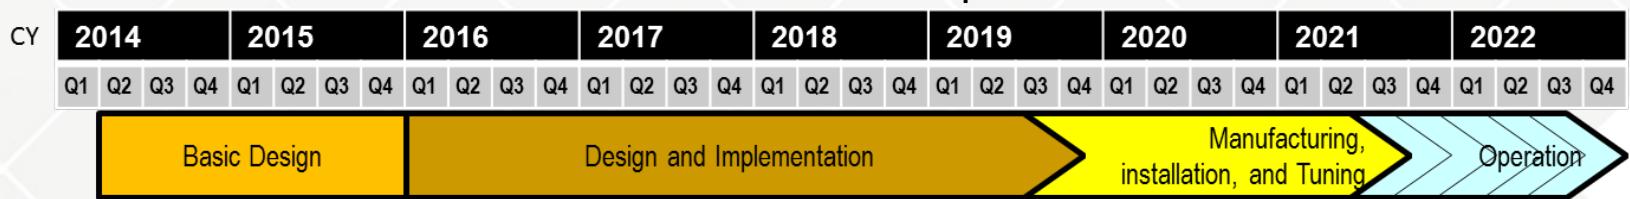
- Building the Japanese national flagship supercomputer, post K, and
- Developing wide range of HPC applications, running on post K, in order to solve social and science issues in Japan

## □ Project organization

- Post K Computer development
  - RIKEN AICS is in charge of development
  - Fujitsu is vendor partner.
  - International collaborations: DOE, JLESC, ..
- Applications
  - The government selected 9 social & scientific priority issues and their R&D organizations.

## □ Status and Update

- “Basic Design” was finalized and now in “Design and Implementation” phase.
  - Now, we are working on detail evaluation by simulators and compilers
- We have decided to choose ARM v8 with SVE as ISA for post-K manycore processor.
- Some delay of delivery will be expected.



# Target Architecture in near future

- We also have two branched projects from EigenExa on the K computer architecture
  - GPU (single):
    - Eigen-G = Experimental code on a single node + a single GPU environment
    - ASPEN.K2 = Automatic-tuning GPU BLAS kernels, especially, SYMV kernel
  - Intel Xeon Phi
    - Divide and conquer algorithm for GEVP focused on a pair of banded matrices



TI, etc, "Eigen-G: GPU-based eigenvalue solver for real-symmetric dense matrices", PPAM2013, LNCS8384  
TI, etc. "High Performance SYMV Kernel on a Fermi-core GPU", VECPAR 2012, LNCS 7851,  
TI, etc. "Automatic-tuning for CUDA-BLAS kernels by Multi-stage d-Spline Pruning Strategy", @^2HPSC 2014



Y.Hirota, etc, "Parallel Divide-and-Conquer Algorithm for Solving Tridiagonal Eigenvalue Problems on Manycore Systems", PPAM2017, LNCS10777  
Y.Hirota, etc. "Acceleration of Divide and Conquer Method for Generalized Eigenvalue Problems of Banded Matrices on Manycore Architectures", PMAA14.

# For Very small kernel

- **Kevd: Developed by Mr. Kudo (was a Ph.D student)**

- For acceleration of a block Jacobi iteration
- Wilkinson's overlap technique (SYR2(K)+SYMV)
- 4x4 micro memory tiling to make use of cache line maximum so called SOSOA
- SIMD intrinsic

- **Extension**

- 8x4 or 8x8 tiling
- AVX2, AVX-512 for KNL

ARM-SVE for post-K

```
! Wilkinson's technique
! Blue: DSYMV
! Red : DSYR2
do i=1,m
    ui=u(i), vi=v(i), yi=y(i), xi=x(i)
    do j=1,i-1
        aij = a(j,i)
        aij += u(j)*vi + v(j)*ui
        y(j) += aij * xi
        yi += aij * x(j)
        a(j,i) = aij
    enddo
enddo
```

# Kevd on K and KNL

- Benchmark and Performance comparison

- K: SPARC64 VIII+fx (by Mr.Kudo and Prof.Yamamoto)
- x86 (Xeon)
- KNL

- Wilkinson512 on NKL

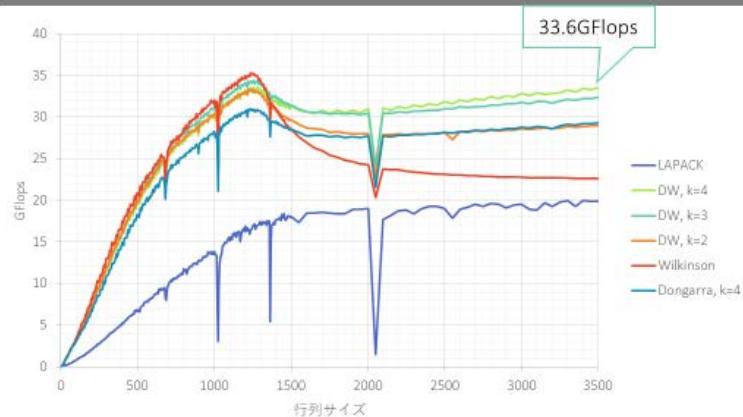
- -O3 -xHost -qopenmp
- -lmkl\_intel\_lp64 -lmkl\_intel\_thread -lmkl\_core -liomp5 -lpthread -lm -ldl

## Rough evaluation

- 2~7GFLOPS/core

## K computer

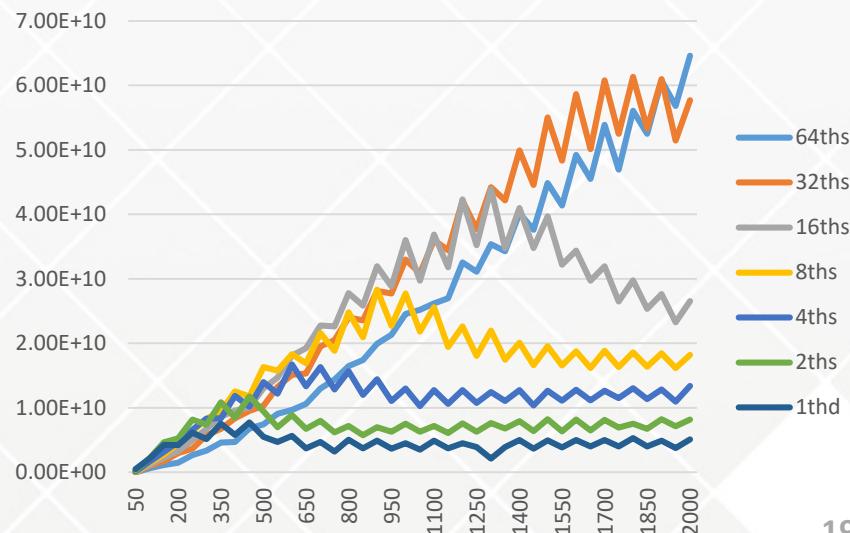
三重対角化の  
実行性能の測定結果



テスト行列：5～1500まで5刻み、1550から3500まで50刻みの大きさの  
minij行列： $A = \{a_{ij}\}, a_{ij} = \min(i, j)$ .

24

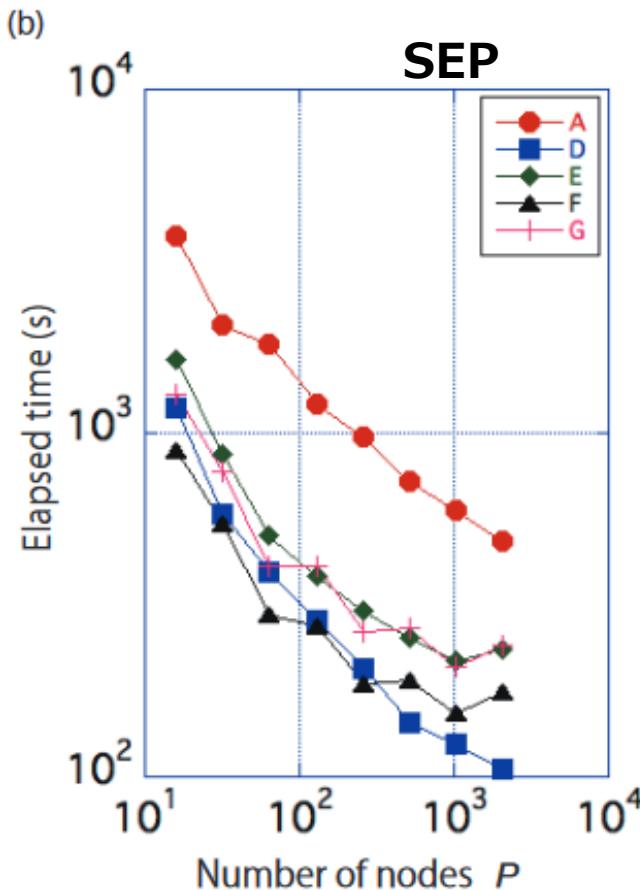
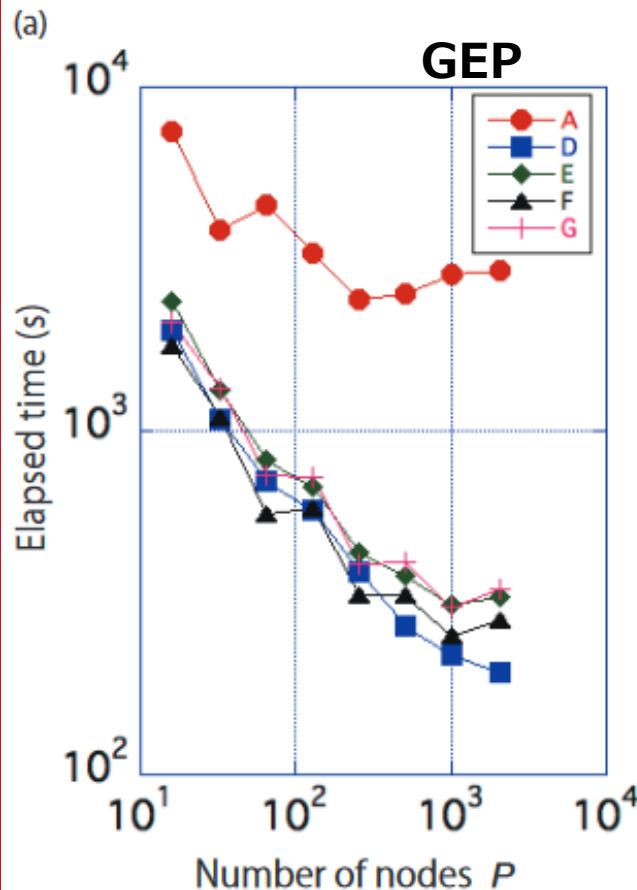
## Performance [FLOPS] on KNL



19

# EigenKernel for GEP (Prof.Hoshi@Tottori U.)

- ELPA 1|2(+Cholesky)+EigenExa on OFP (a KNL cluster)



Workflow	SEP solver	Reducer
A	ScalAPACK	ScalAPACK
B	Eigen_sx	ScalAPACK
C	ScalAPACK	ELPA
D	ELPA2	ELPA
E	ELPA1	ELPA
F	Eigen_s	ELPA
G	Eigen_sx	ELPA

**Fig. 3** Benchmark on Oakforest-PACS. The matrix size of the problem is  $M = 90,000$ . The computation was carried out with  $P = 16, 32, 64, 128, 256, 512, 1024, 2048$  nodes in the workflows of  $A$ (circle),  $D$ (square),  $E$ (diamond),  $F$ (triangle),  $G$ (cross). (a) The elapsed time for the whole GEP solver. (b) The elapsed time for the reduced SEP solver.

# Agenda

## 1. Quick Overview of Project

- Past and Present of EigenExa
- Diagonalization of a 1million x 1million matrix on K computer

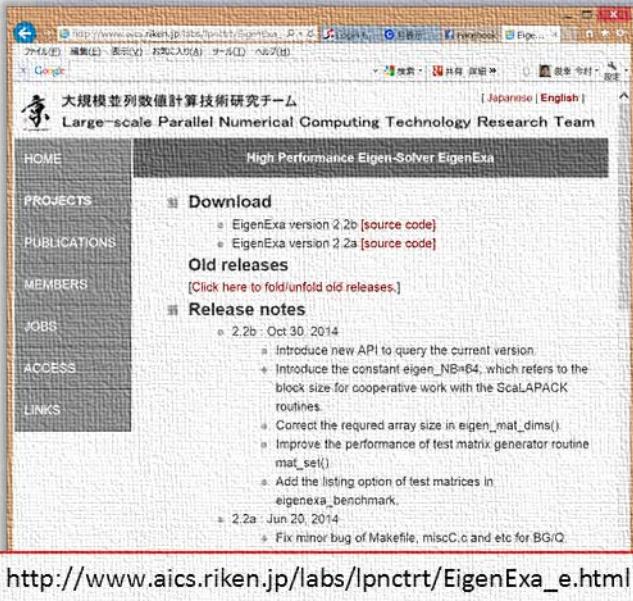
## 2. The latest updates

## 3. Future direction

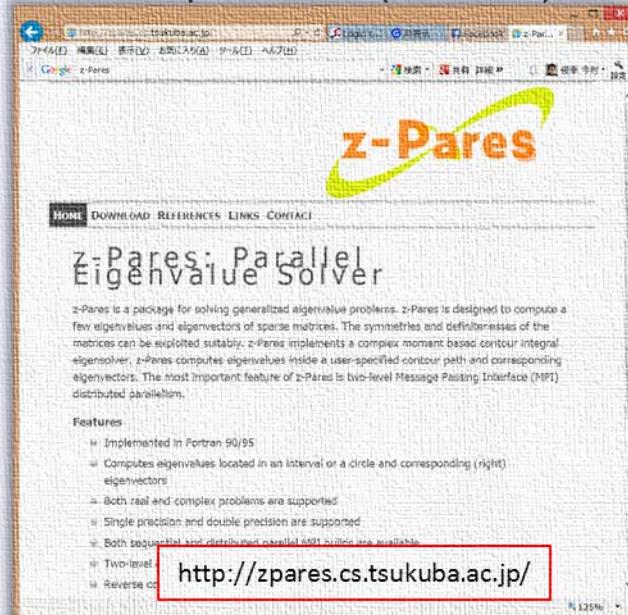
## 4. Summary

### Prof. Sakurai Team 'H4ES'

EigenExa:: dense solver (RIKEN)



Z-Pares:: sparse solver (U.Tsukuba)



# Summary of talk

- **EigenExa project (2011-2016, 2017-)**
  - The first milestone : 1million order eigenvalue computation with full nodes of K computer.
  - Second milestone : optimization of communication
  - Current version 2.4 b is available from
    - ✓ <http://www.r-ccs.riken.jp/labs/lpnctr/en/projects/eigenexa/>
    - ✓ [http://www.r-ccs.riken.jp/labs/lpnctr/assets/img/EigenExa-2.3c\\_users\\_manual\\_2014-06-24\\_en.pdf](http://www.r-ccs.riken.jp/labs/lpnctr/assets/img/EigenExa-2.3c_users_manual_2014-06-24_en.pdf)
- **Future work towards post-K project also exa-scale computing**
  - Establish the CA technology for total performance of EigenExa
  - Reduced/Quadruple Precision version
  - Vector computers, other platforms
- **Our goal is reality on Exa-scale computing!**
  - New target architecture, ARM64+SVE
  - New topic is also concerned like Reproducibility

THANKS!