Performance benchmark of the latest EigenExa on Fugaku

Toshiyuki Imamura
Takeshi Terao
Takuya Ina*
imamura.toshiyuki@riken.jp
takeshi.terao@riken.jp
RIKEN R-CCS
Kobe, Hyogo, Japan

Yusuke Hirota
University of Fukui
Fukui-shi, Fukui, Japan
y-hirota@u-fukui.ac.jp

Katsuhisa Ozaki
Yuki Uchino
ozaki@sic.shibaura-it.ac.jp
Shibaura Institute of Technology
Saitama-shi, Saitama, Japan

1 INTRODUCTION

Eigensolver is one of the strongly demanded tools in modern simulations. RIKEN has been developing high-performance and reliable eigensolvers since 2012 for the K computer, 2020 for the supercomputer Fugaku. 'EigenExa' [3] refers to the heart of our eigensolver project and the signature software for large-scale and highly parallel computations. We have just released version 2.11 in 2021 as a standard library on Fugaku, which employs the novel one-stage scheme. Also, the traditional one-stage scheme is available as most of the state-of-the-arts eigensolvers such as ScaLAPACK and ELPA.

2 EIGENEXA

The latest RC version of EigenExa v2.12 supports (only FP64);

- Real-symmetric Standard eigenvalue problem,
- Real-symmetric Generalized eigenvalue problem,
- Hermitian Standard eigenvalue problem, and
- Hermitian Generalized eigenvalue problem.

Along with investigations for modern parallel systems, a band-matrix version of the one-stage scheme, communication avoiding algorithms, and kd-tree based load balancing are introduced.

The difficulties to eliminate the nondeterminism of parallel computation and to ensure perfect reproducibility are common issues in the HPC field. EigenExa is coded elaborately to guarantee as much reproducible as possible and achieves high performance, while ScaLAPACK is not permanently reproducible for consecutive calls.

Like most eigensolvers, EigenExa exhibits typical accuracy behavior, dominated by the condition number and the gap of the eigenvalues. The idea to refine the accuracy is another approach, and we are currently attempting to include an accuracy booster according to Ogita-Aishima's [2] scheme into version 2.12 or later.

Numerical verification is also investigated with a new approach using only floating-point arithmetic with nearest-point rounding [4]. The method's highlights are high-precision matrix multiplication to cope with large matrices and to suppress the overestimation of the error upper bound even for ill-conditioned matrices.

3 BENCHMARK ON FUGAKU

The page regulation allows us to show only the preliminary performance comparison among `P{DSY|ZHE}EVD` of ScaLAPACK and `eigen_{FS|h}` of EigenExa in Figure 1, which perform the real symmetric/Hermitian standard eigenproblem for all spectrum. The benchmark was done on Fugaku housed at RIKEN using up to 16,384 nodes (2.2GHz boost mode). Observations reveal greater parallel scaling of EigenExa, with a 3-8x performance improvement over ScaLAPACK in the real symmetric case. The performance profile of Hermitian on EigenExa exhibits a similar trend (1.5-4x improvement) to the real symmetric case. However, `PDSYEVD` and `PZHEEVD` show almost equivalent or upset performance, which is rather unnatural in terms of 2x memory traffic and 2x computational intensity between real and complex, thus to be clarified.

For $N = 10^6$, a remarkable benchmark of 4,096 nodes has recorded 1.37 PFLOPS in 3,100 seconds, which is comparable or superior to those ever done on the K computer and a Fujitsu FX10 [1].

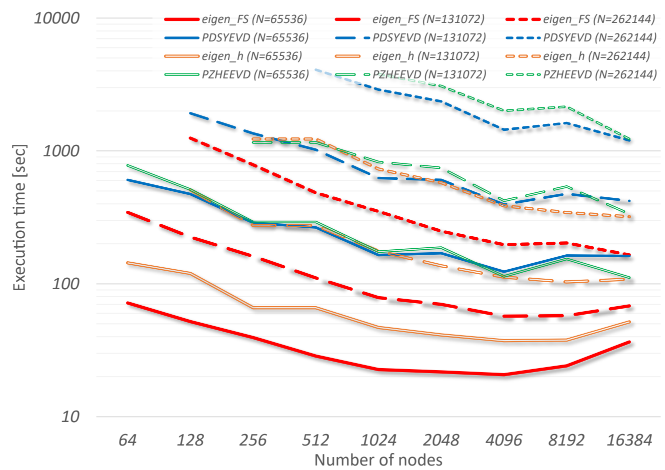

Figure 1: Performance comparison `P{DSY|ZHE}EVD` versus `eigen_{FS|h}` on Fugaku with 1ppn configuration

ACKNOWLEDGMENTS

The research was supported by the Fugaku trial phase program, and used computational resources of the supercomputer Fugaku provided by RIKEN R-CCS (ra000006, hp200263, and hp200313).

REFERENCES

[1] H. Imachi and T. Hoshi. 2016. Hybrid Numerical Solvers for Massively Parallel Eigenvalue Computations and Their Benchmark with Electronic Structure Calculations. *Journal of Information Processing* 24, 1 (2016), 164–172.
[2] T. Ogita and K. Aishima. 2018. Iterative Refinement for Symmetric Eigenvalue Decomposition. *Japan Journal of Industrial and Applied Mathematics* 35, 3 (2018), 1007–1035.
[3] RIKEN R-CCS 2021. EigenExa homepage. https://www.r-ccs.riken.jp/labs/lpnctrt/projects/eigenexa/.
[4] T. Terao, K. Ozaki, T. Ogita, and T. Imamura. 2021. Verified Numerical Computations for a Standard Eigenvalue Problem without Directed Rounding. The 40th JSST Annual International Conference on Simulation Technology.

*Currently, Japan Atomic Energy Agency