Large-Scale First-Principles Electronic Structure Calculations in Petascale and Exascale Supercomputers: A Real-Space Density Functional Theory code

J.-I. Iwata
Department of Applied Physics, The University of Tokyo

http://github.com/j-iwata/RSDFT

RSDFT
実空間密度汎関数法

Interdisciplinary symposium on modern density functional theory
Jun.19 - Jun.23, 2017

Okochi Hall, RIKEN

CdMSi

科研費
KAKENHI

FIRST-PRINCIPLES ELECTRONIC STRUCTURE CALCULATIONS SPREAD IN A WIDE RANGE OF FIELDS WITH SOPHISTICATED PROGRAM PACKAGES

- Density Functional Theory
  - relatively low computational costs : $O(N^3)$
  - reasonable accuracy
- Used in the R&D in industry
  - realistic materials → large-scale systems need to be handled
- Many program packages
  - VASP, QUANTUM ESPRESSO, ABINIT, CPMD, PHASE, STATE, xTAPP, QMAS, CASTEP, DACAPO, etc.
    (listed up only the plane-wave pseudopotential programs here)

VARIOUS METHODS OTHER THAN THE PLANE-WAVE METHOD

- Other basis functions and representative packages
  - atomic orbital (SIESTA)
  - pseudo atomic orbital (OpenMX)
  - wavelet (BigDFT)
  - Finite-element (FEMTECK)
  - Finite-difference (PARSEC, GPAW, RSDFT)

- Drawbacks of the plane-wave method
  - Periodic Boundary Condition
    - super-cell models must be used for 0-D～2-D systems
  - Delocalized in real space
    - Incompatible with the order-N method
  - Fast Fourier Transform
    - communication burden on massively-parallel computing

Top seven supercomputers for Nov 2016

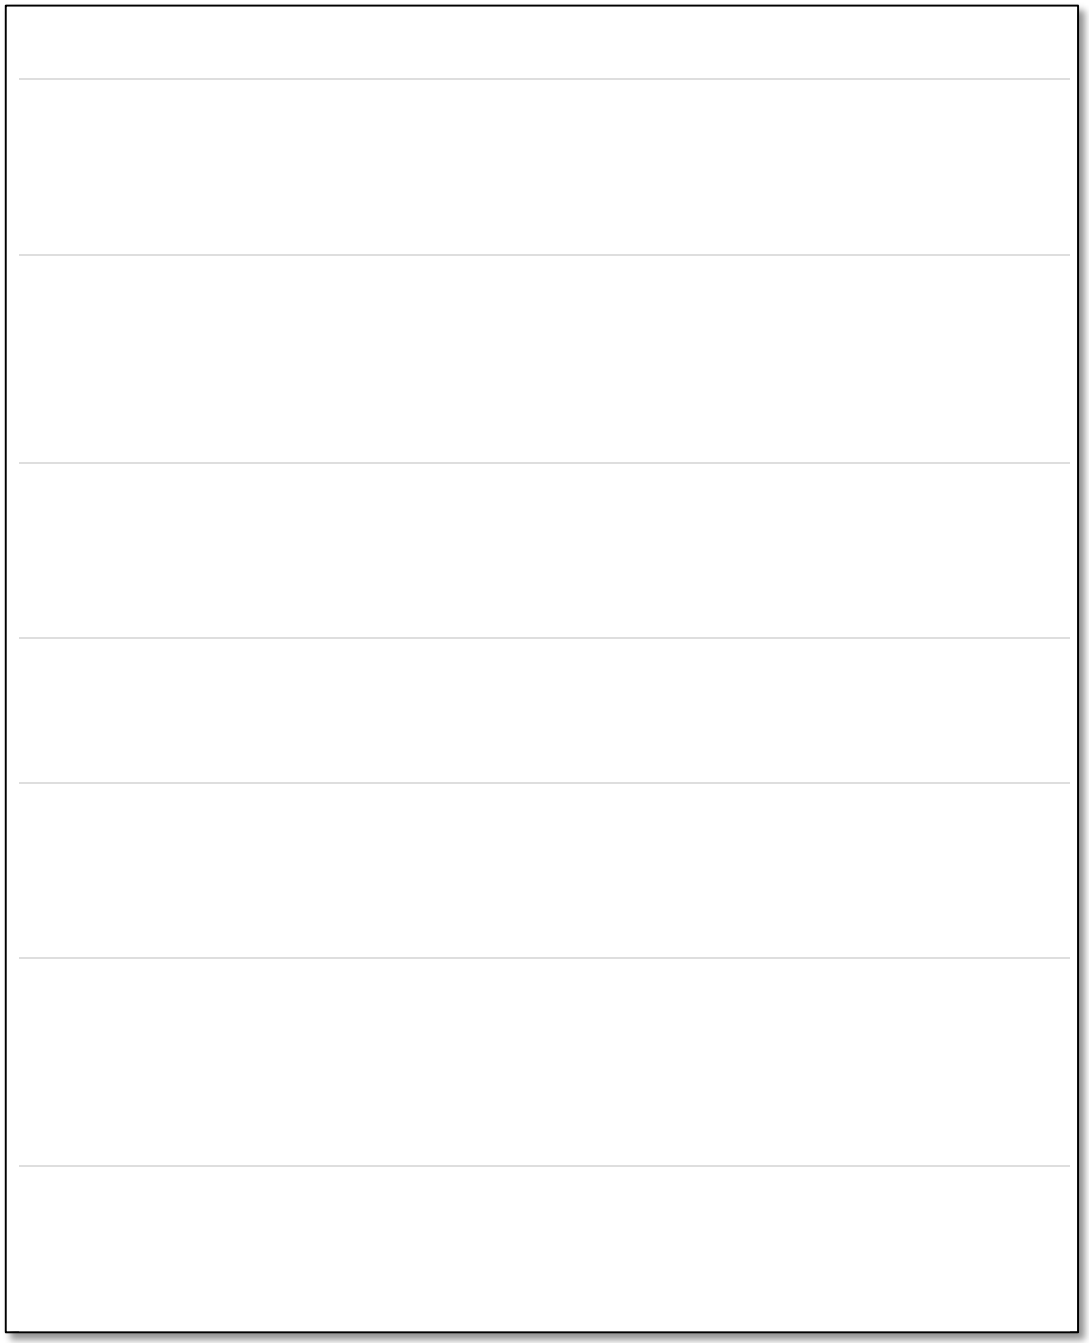

https://www.top500.org/

FROM K COMPUTER TO POST-K COMPUTER

K computer
SPARC64 VIIIfx
8 cores/CPU (128 GFLOPS)

FX100
SPARC64 XIfx
32 + 2 cores/CPU (>1 TFLOPS)

Oakforest-PACS
Intel Xeon Phi 7250 (Knights Landing)
68 cores/CPU (3.0464 TFLOPS)
1 CPU/node
8208 nodes (25 PFLOPS)

POST-K COMPUTER

- The next flagship machine is being developed (～2021)
- RIKEN AICS and FUJITSU
- ARM architecture
- Multi core processor ➔ Many core processor
  - Core-Memory Group (CMG)
    - several cores are grouped together
    - MPI processes run on each CMG

REAL-SPACE FINITE-DIFFERENCE PSEUDOPOTENTIAL METHOD

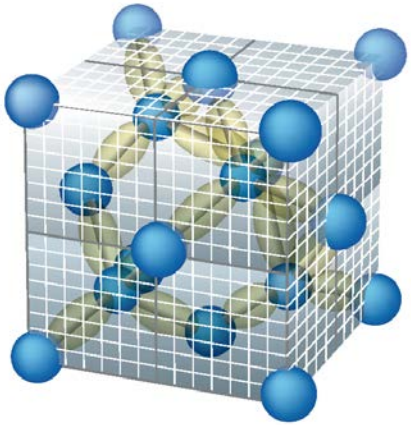

$$\begin{cases} \left(-\frac{1}{2}\nabla^2 + \hat{v}_{ion}^{PP} + \int d\mathbf{r}' \frac{\rho(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|} + \frac{\delta E_{XC}[\rho]}{\delta\rho(\mathbf{r})}\right)\phi_i(\mathbf{r}) = \varepsilon_i\phi_i(\mathbf{r}) \\ \rho(\mathbf{r}) = \sum_{i=1}^{N}|\phi_i(\mathbf{r})|^2 \qquad (N << \#\text{ of grid points}) \end{cases}$$

Kohn-Sham equation is solved as a finite-difference equation in discretized space

- **Derivatives** → (higher-order) finite difference $\frac{\partial^2}{\partial x^2}\phi(\mathbf{r}) \approx \sum_{m=-6}^{6} C_m \phi(x+m\Delta, y, z)$
- **Integrals** → summation over grid points $\int \phi_m^*(\mathbf{r})\phi_n(\mathbf{r})d\mathbf{r} \approx \sum_{i=1}^{N_{grid}} \phi_m^*(\mathbf{r}_i)\phi_n(\mathbf{r}_i)\Delta V$
- **Ionic potentials** → Pseudopotentials $\hat{v}_{ion}^{PP} = v_{local}^{PP}(\mathbf{r}) + \sum_{a,l,m} |\beta_{alm}\rangle\langle\beta_{alm}|$

N. Troullier & J. L. Martins, Phys. Rev. B 34, 1993 (1991)

HISTORY OF RSDFT

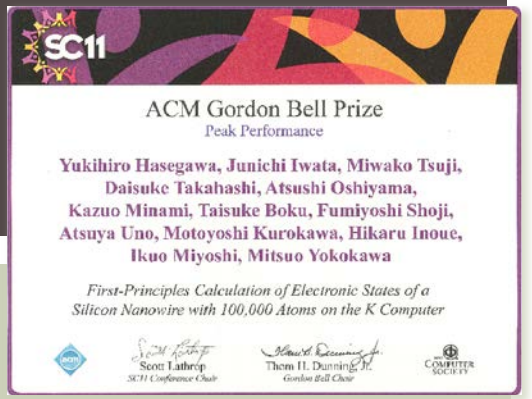
SC11

ACM Gordon Bell Prize

Peak Performance

Yukihiro Hasegawa, Junichi Iwata, Miwako Tsuji,
Daisuke Takahashi, Atsushi Oshiyama,
Kazuo Minami, Taisuke Boku, Fumiyoshi Shoji,
Atsuya Uno, Motoyoshi Kurokawa, Hikaru Inoue,
Ikuo Miyoshi, Mitsuo Yokokawa

First-Principles Calculation of Electronic States of a Silicon Nanowire with 100,000 Atoms on the K Computer

Scott Lathrop, SC11 Conference Chair

Thom H. Dunning, Jr., Gordon Bell Chair

COMPUTER SOCIETY

Three-dimensional time-dependent Hartree-Fock calculation : Application to $^{16}O + ^{16}O$ collisions
H. Flocard, S. E. Koonin and M. S. Weiss
Phys. Rev. **C17**, 1682 (1978).

Higher-order finite-difference pseudopotential methodd: An application to diatomic molecules
J. R. Chelikowsky, N. Troullier, K. Wu, and Y. Saad
Phys. Rev. **B50**, 11355 (1994).

Time-dependent local-density approximation in real time
K. Yabana and G. F. Bertsch
Phys. Rev. **B54**, 4484 (1996).

Real-space, real-time method for the dielectric function
G. F. Bertsch, J.-I. Iwata, A. Rubio, and K. Yabana
Phys. Rev. **B62**, 7998 (2000).

RSDFT
実空間密度汎関数法

J.-I. Iwata *et al.*, J. Comp. Phys. 229, 2339 (2010).
http://github.com/j-iwata/RSDFT

AWARD
ACM Gordon Bell Prize (2011)

Y. Hasegawa, *et al.*, SC'11 Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis

PARALLELIZATION IN RSDFT

Real Space Grids

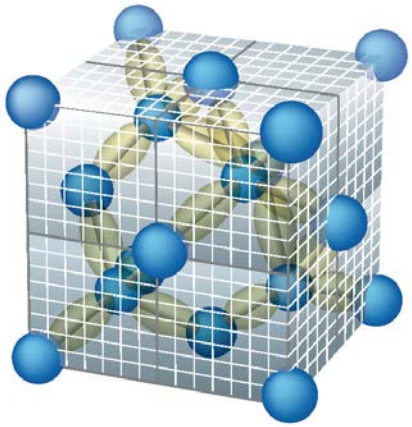

CPU Space

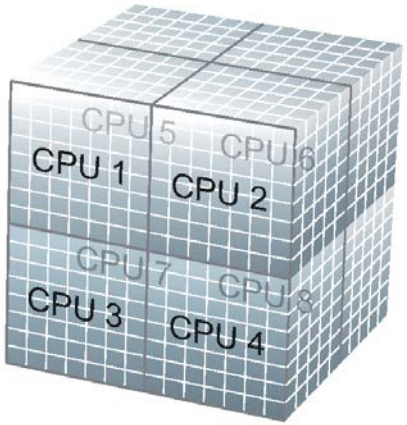

J.-I. Iwata *et al.*, J. Comp. Phys. 229, 2339 (2010).

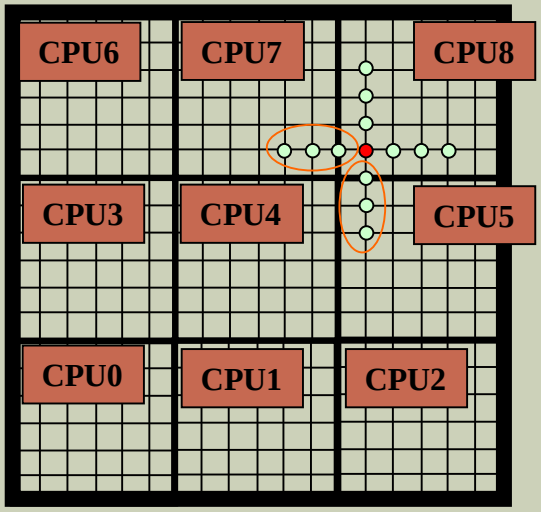

- MPI (Message-Passing Interface) library
  MPI_ISEND, MPI_IRECV → finite-difference calc.
  MPI_ALLREDUCE → global summation
- OpenMP
  Further grid parallelization (within each node, CPU, or CMG) is performed by thread parallelization
- spin, k-point, and band index parallelization (orbital parallelization) is also implemented

FLOW CHART OF THE STATIC DFT CALCULATION

Subspace-iteration method for large-sparse eigenvalue problem

1 (**CG**) Conjugate-Gradient method

$$H=\left(-\frac{1}{2}\nabla^2+\hat{v}_{ion}^{PP}+\int d\mathbf{r}'\frac{\rho(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|}+\frac{\delta E_{XC}[\rho]}{\delta\rho(\mathbf{r})}\right)$$

$$\varepsilon_n=\frac{\langle\phi_n|H|\phi_n\rangle}{\langle\phi_n|\phi_n\rangle}$$

→ minimize

O(N²)

2 (**GS**) Ortho-normalization by Gram-Schmidt method

O(N³)

3 Density and potential update

O(N)

4 (**SD**) Subspace Diagonalization

Diagonalize H within a small dimensional space

spanned by $\{\phi_1(\mathbf{r}),\phi_2(\mathbf{r}),\cdots,\phi_N(\mathbf{r})\}$

$$h_{mn}=\langle\phi_m|H|\phi_n\rangle$$

O(N³)

Almost all $O(N^3)$ computations can be performed with LEVEL 3 BLAS

Gram-Schmidt procedure

$$\psi_1 = \varphi_1$$

$$\psi_2 = \varphi_2 - \psi_1\langle\psi_1|\varphi_2\rangle$$

$$\psi_3 = \varphi_3 - \psi_1\langle\psi_1|\varphi_3\rangle - \psi_2\langle\psi_2|\varphi_3\rangle$$

$$\psi_4 = \varphi_4 - \psi_1\langle\psi_1|\varphi_4\rangle - \psi_2\langle\psi_2|\varphi_4\rangle - \psi_3\langle\psi_3|\varphi_4\rangle$$

$$\psi_5 = \varphi_5 - \psi_1\langle\psi_1|\varphi_5\rangle - \psi_2\langle\psi_2|\varphi_5\rangle - \psi_3\langle\psi_3|\varphi_5\rangle - \psi_4\langle\psi_4|\varphi_5\rangle$$

$$\psi_6 = \varphi_6 - \psi_1\langle\psi_1|\varphi_6\rangle - \psi_2\langle\psi_2|\varphi_6\rangle - \psi_3\langle\psi_3|\varphi_6\rangle - \psi_4\langle\psi_4|\varphi_6\rangle - \psi_5\langle\psi_5|\varphi_6\rangle$$

$$\psi_7 = \varphi_7 - \psi_1\langle\psi_1|\varphi_7\rangle - \psi_2\langle\psi_2|\varphi_7\rangle - \psi_3\langle\psi_3|\varphi_7\rangle - \psi_4\langle\psi_4|\varphi_7\rangle - \psi_5\langle\psi_5|\varphi_7\rangle - \psi_6\langle\psi_6|\varphi_7\rangle$$

$$\psi_8 = \varphi_8 - \psi_1\langle\psi_1|\varphi_8\rangle - \psi_2\langle\psi_2|\varphi_8\rangle - \psi_3\langle\psi_3|\varphi_8\rangle - \psi_4\langle\psi_4|\varphi_8\rangle - \psi_5\langle\psi_5|\varphi_8\rangle - \psi_6\langle\psi_6|\varphi_8\rangle - \psi_7\langle\psi_7|\varphi_8\rangle$$

Matrix × Matrix product

Recursive algorithm

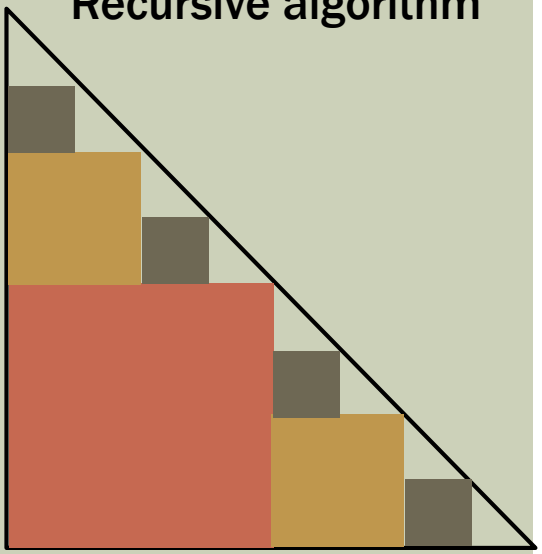

BENCHMARK TEST @K : 100,000-ATOM Si NANOWIRE

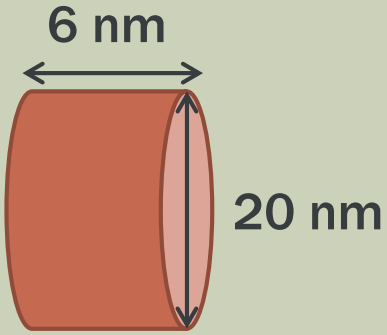

(110) SiNW

# of atoms :	107,292
# of grid points :	576 x 576 x 192 = 63,700,992
# of orbitals :	229,824

# of CPUs : 55,296 (442,368 cores)
(18,432 (grid para.) x 3 (orbital-para.))

heoretical Peak : 7.07 PFLOPS

Computational time of one step of SCF iteration

	Execution time (sec.)	Computation time (sec.)	Performance (PFLOPS)
SCF	5456	4417	3.08 (43.6 %)
DIAG	3710	3218	2.72 (38.5 %)
CG	209.3	57.66	0.05 (0.7 %)
GS	1537	1141	4.37 (61.8 %)

FULL-NODE TEST @K

system : Si nanowire (107,292 atoms)
grid : 576 x 576 x 180 = 59,719,680
band : 230,400

nodes: 82,944 (10.6 PFLOPS
cores : 663,552

# of atoms = 39.696
Grid = 348 x 348 x 288
Orbitals = 82,176
# of CPUs = 32 x 22 x 6 = 4224 (541 TFLOPS)
38 steps for SCF convergence → 45 hours

	Execution time (sec.)
SCF	3837
SD	2171
CG	308.5
GS	1313

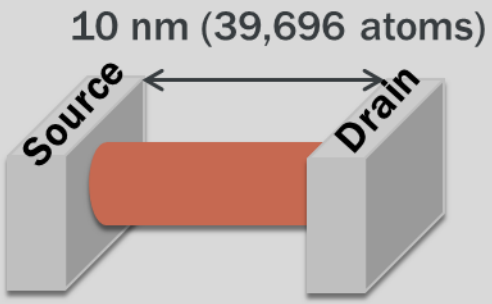

Elapsed time for 1 step of S

	SCF	CG O(N²)					
time (s)	2931	160					
PFLOPS	5.48	0.06					
peak%	51.67	0.60	63.10	50.17	57.93	1.03	76.70

Y. Hasegawa *et al.*, J. High Perform. Comp. Appl. 28, 335 (2014)

TWISTED BILAYER GRAPHENE

K. Uchida, S. Furuya, JI, A. Osiyama, Phys. Rev. B (2014)

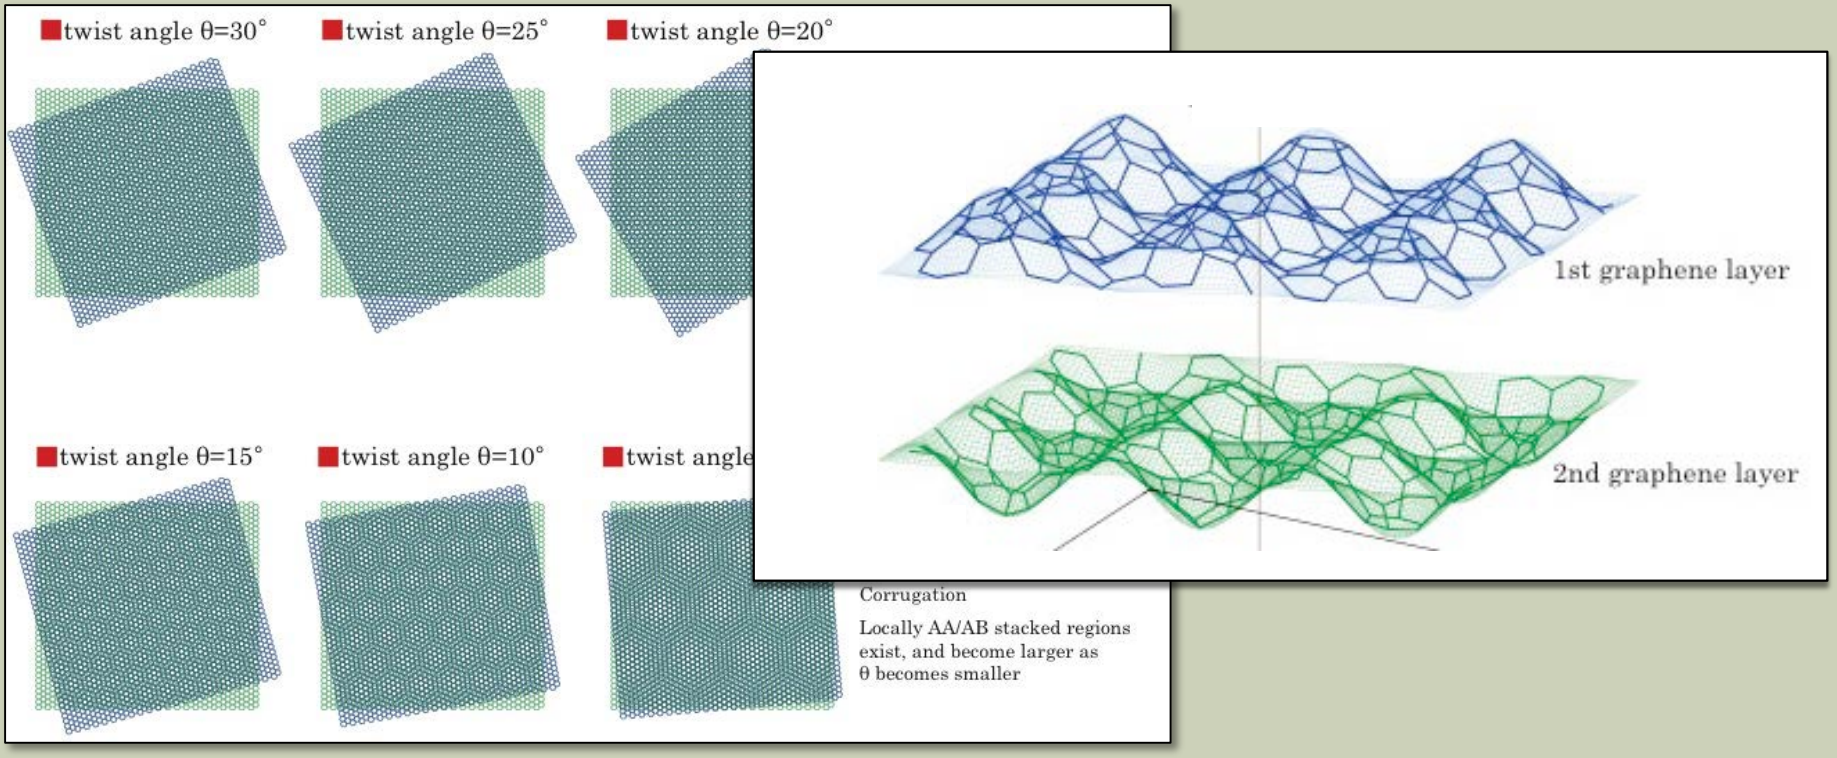

Elapsed time for 1 step of SCF loop at K computer (841 nodes)

atoms	SCF	CG	GS	SD
22,708	2980	446	710	1452

(sec.)

Fermi-velocity Reduction in tBLG

Poster no.11
Hirofumi Nishi

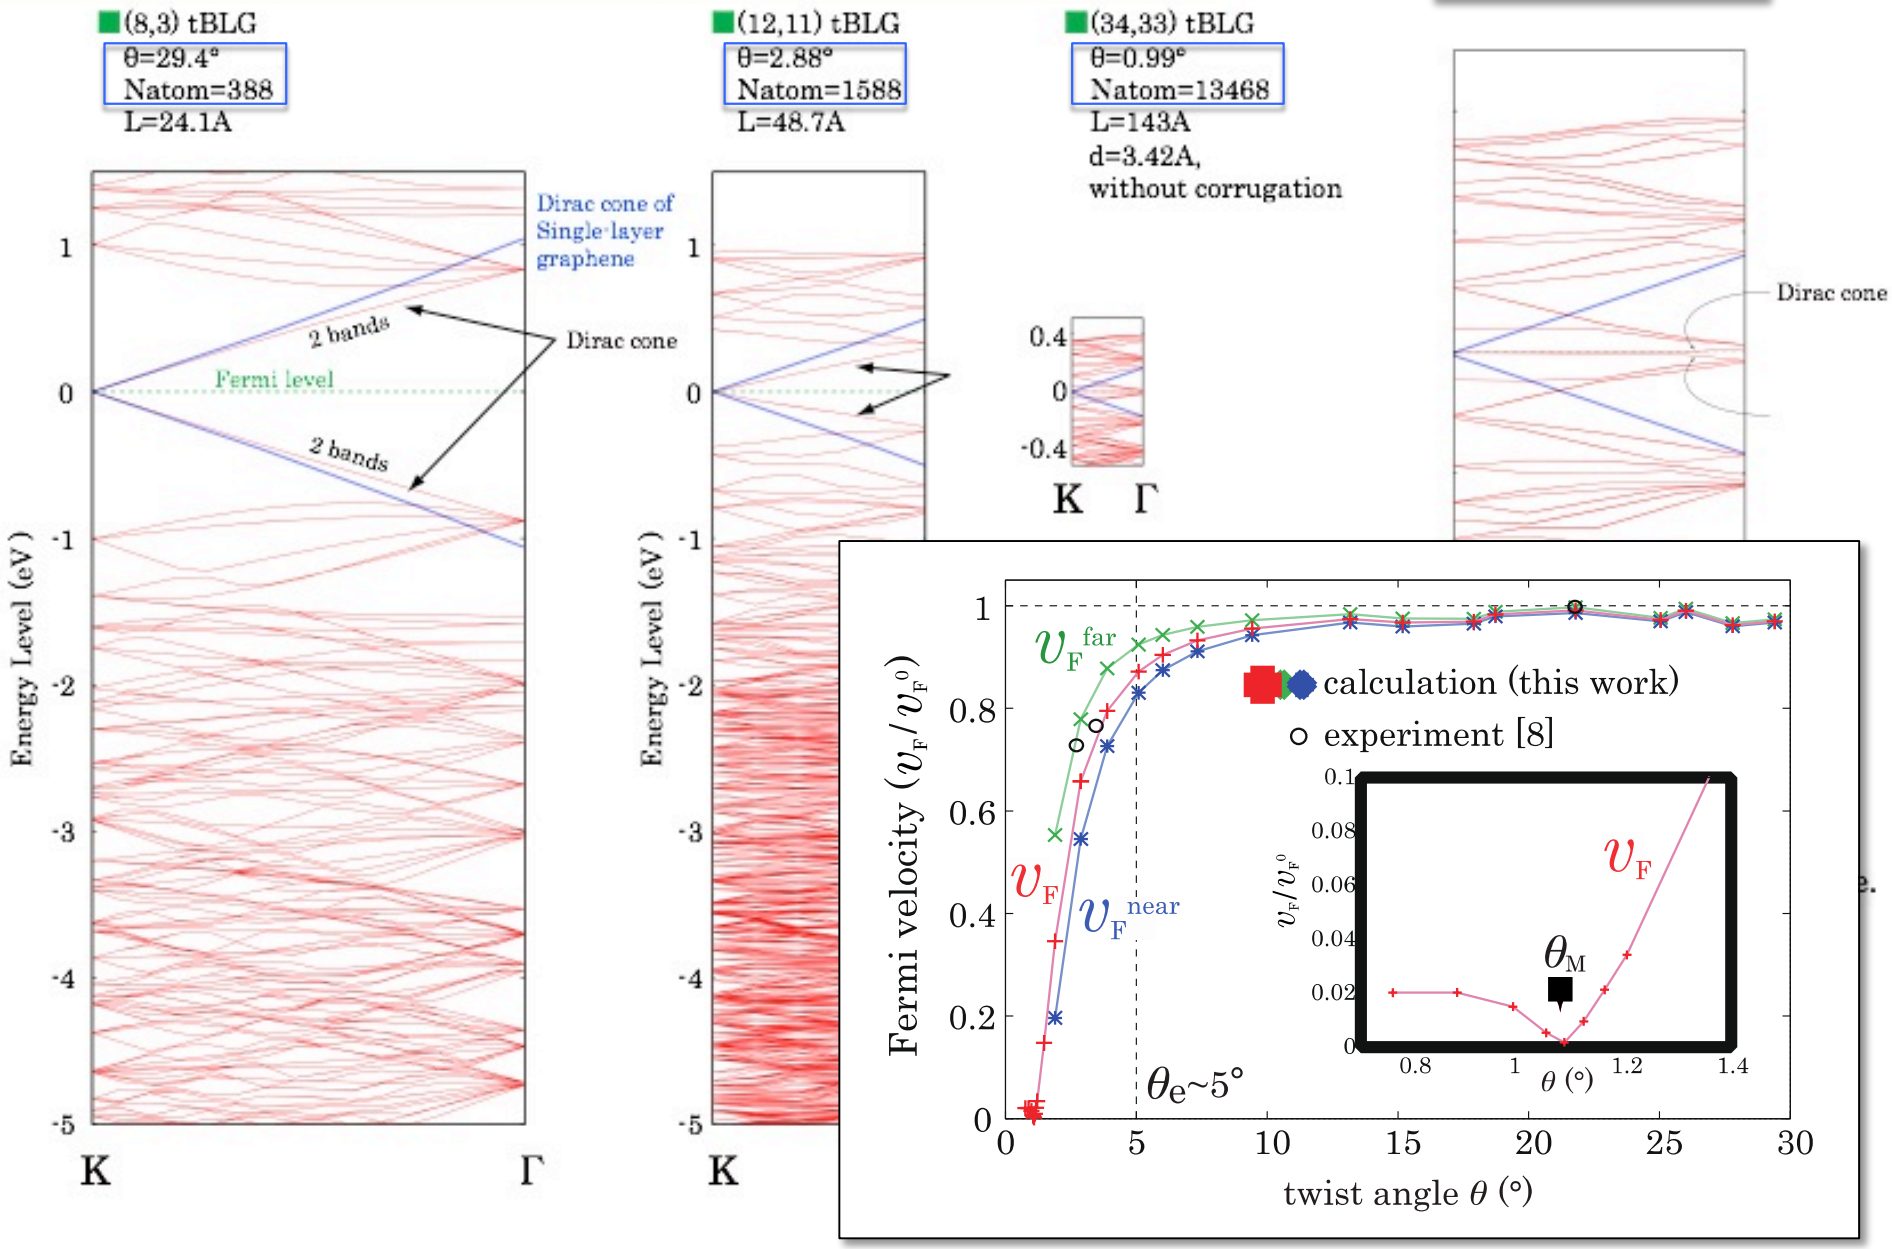

SUMMARY

http://github.com/j-iwata/RSDFT

- RSDFT
  - A first-principles electronic structure calculation code based on the real-space finite-difference pseudopotential method
  - More suitable for massively-parallel computing than that of PW basis
- K computer (computable systems)
  - several thousand atoms with atomic structure opt. (～ 4000)
  - a few tens of thousand atoms for single-point calc. (～ 40,000)
- post-K computer (～ 1 EFLOPS)
  - tens of thousands of atoms with atomic structure opt.
  - the system, which is feasible only with the whole resource of the K computer, is expected to be computable with a small resource of the post-K computer