Performance Evolution of the FLEUR Code

U.Alekseeva, G.Michalicek and D. Wortmann

Peter Grünberg Institut and Institute for Advanced Simulation. Forschungszentrum Jülich and JARA, 52425 Jülich, Germany

Abstract

The advances in performance optimisation of the all-electron DFT code FLEUR achieved during the MaX project are presented. The improved implementation features increased modularity of the code, reduced I/O, hybrid MPI/OpenMP parallelisation, additional interfaces to external libraries providing performance portability and provides a significantly improved users experience. Due to the performance boost of the MaX version of the FLEUR code simulations with unit cells of more than 1000 atoms are now feasible

Unique Set of Features

www.flapw.de

eur

At PGI-1/IAS-1 we develop and use the advanced DFT-code FLEUR: Highly accurate FLAPW method

- All electron, full-potential DFT
- Bulk systems, surfaces, molecules
- Complex magnetism, including non-collinear magnetism.
- calculation of model parameters for localized spin models Relativistic effects, spin-orbit coupling
- Applied external electric fields
- Standard DFT exchange and correlation functionals as well
- as LDA/GGA+U, hybrid functionals, vdW functionals Interfaces to external codes like Wannier90, SPEX, GFLEUR

Algorithms for FLAPW

- A standard FLEUR run performs the following main tasks: 1. The calculation of the potentials from the charge density
- Coulomb potential as solution of Poisson equation
 - Exchange-correlation potential
- 2. The setup of the Hamiltonian and overlap matrices Interstitial contribution in plane-waves

 - Sphere contributions in local basis Matching of local basis to plane-waves
- 3. The diagonalization of the matrix
- Call to optimized math-libraries
- 4. Calculation of new density
 - Sphere and interstitial density Mixing of in- and output density

```
FLAPW-basis set defined in two different regions in space:
```

- $\begin{array}{l} \text{Plane-waves} \\ \text{Numerical (atom-like)} \\ \text{functions in MT-sphere} \end{array} \phi_{j}^{\text{LAPW}} = \begin{cases} e^{i(\vec{k}+\vec{g}_{j})\vec{r}} \\ \sum_{lm}^{L} A_{j,lm} \phi_{lm}^{\text{atomic}}(\vec{r}) \end{cases} \end{cases}$
- around each atom Matching across sphere boundary (A-coefficients)

Sphere contribution to Hamiltonian:



Optimized Algorithm: Separation into spherical and non-spherical contribution $H_{i,j}^{\mathrm{MT}} = \sum \ H_{i,j}^{\mathrm{loc;sph}} + H_{i,j}^{\mathrm{loc;non-sph}}$

Spherical contribution:

m,m'-sum can be performed analytically

remaining single I-sum is fast $H_{i,j}^{\rm loc;sph} = \sum a_{i,l}^* H_{l,l}^{\rm loc;sph} a_{j,l} D_{i,j}$







(Itwill Addi	en oniversity)	JAF	RAHPC				
Identification of computational most relevant code segments							
Code part	CPU time (24 core)	Speedup on 384 cores	Parallel efficiency				
Setup,Potential	5%	3.1	20%				
Matrix setup	45%	10	63%				
Diagonalization	40%	9.0	56%				
Charge density	8%	5.3	33%				
Mixing	1%	2.2	14%				
Total SCF run	100%	7.8	49%				
256 atoms Cu	Ag on CLAIX,	scales up to 16 n	odes (384 cores)				

Modernization of the Code

Intensive support and discussions with HPC groups at JSC

(Jülich Forschungszentrum) and IT Center

(RWTH Aachen University)

Implementation of multi-level hybrid parallelism enabling the efficient use of multi-core and multi-node machines using MPI and OpenMP

Code Parts	Level of Parallelization					
	MPI		OpenMP	SIMD		
Potential	 a lot of small subroutines parallelization of loops 		- loops	 compiler flags BLAS calls 		
k – points Eigenvalue						
Matrix Setup	 independent eigenvalue problems for each 	- block distribution among the processes	- BLAS calls - compiler hints - compiler flags			
Diagonalization	k-point	- interfaces to the external libraries: ELPA, ScaLAPACK, Elemental				
New Charge	- loop over reciprocal lattice vector g		 loop over atoms 	- compiler flag		
Mixing	- to be done					
Version	0.26	Version 0.27 MaX Re	elease 2.0			

- Interfacing of various high-performance libraries for matrix diagonalization (ELPA, ScaLAPACK, Magma, Elemental) Implementation of new algorithms suited for current and
- future computing architectures
- Removal of show-stoppers like excessive IO Code version to use CUDA programming
- for running on GPUs

Improved FLEUR experience

Typical challenges for FLEUR users:

- All-electron method requires many parameters and has a
- complex input
- Input/Output difficult to understand and to modify Availability of documentation

Exemplary actions implemented: < xml/>XML input/output



Tutorials and hands-on

CUDA.

sessions <u>.</u>

Code refactoring and developers tools:

- New version with increased degree of modularization Removing of IO and implementing of structured parallel IO schemes where needed
- Automatic building and testing of executables in a continuous integration framework
- Introduction of modern Fortran programming feature
- Code available for collaborative development on institutes GitLab server:

https://iffgit.fz-juelich.de/fleur/fleur GitLab

Acknowledgment

This work has been supported by the EU through the H2020-EINFRA-2015-1 project: GA 676598 and by a JARA-HPC seed fund project







- Further porting to accelerators
- Further tuning of the KNL performance Redesign of matrix layout to enable easier implementation of
- new algorithms and easier tuning to new architectures