

Estimating the Performance of *Ab Initio* Calculation by VASP on Openpower High Performance System

Vyacheslav E. Lozhnikov¹, Alexander V. Mamonov¹, Vadim O. Borzilov¹, Marina V. Mamonova¹, Pavel V. Prudnikov¹, Aleksei A. Sorokin², Georgy G. Baksheev³

¹ Department of Theoretical Physics, Omsk State University, Omsk, Russia, prudnikovpv@omsu.ru

² Computing Center of Far-Eastern Branch, Russian Academy of Sciences, Khabarovsk, Russia, alsor@febras.net

³ Novosibirsk State University, Novosibirsk, Russia, g.baksheev@gsu.ru

Abstract

In this work we compare the performance of Pascal P100 GPUs vs POWER8 CPU on OpenPOWER HPC system by VASP calculation of energy and magnetic characteristics of Fe/Cu(111)/Fe and Co/Cu(100)/Co multilayer magnetic nanostructures. We revealed that the VASP code demonstrates a maximum performance on OpenPOWER System with the GPUs.

1 Introduction

The behavior of multilayer magnetic structures has become of great technological importance due to the applications in magnetic storage devices. The *ab initio* calculations are widely used to calculate some characteristics of solids [1] and multilayer magnetic structures [2]. The main advantage of *ab initio* approach is independence from experimental data. Unlike the case of semi-empirical methods, there is no need for calibration or fitting parameters. Thus, *ab initio* methods can also be used to calculate the characteristics of perspective systems, i.e., for prediction of properties of materials that have not yet been developed. Mainly used packages that can perform *ab initio* calculations are VASP [3–5], Quantum Espresso [6], ABINIT [7], Wien2K [8].

Effective application of *ab initio* calculations requires the scalability of the code for novel high performance systems (HPS) with different hardware architectures. In this work we focused on the Vienna *Ab initio* Simulation Package (VASP). VASP is a complex package for performing *ab initio* quantum-mechanical molecular dynamics simulations using pseudopotentials or the projector-augmented wave method and a plane wave basis set [4]. Now it is one of the most popular parallel code for quantum chemistry and solid-state calculations of electronic structure. So the estimation of VASP code performance for HPS is an actual and non-trivial task. In this paper, we compare time of execution on POWER8 CPU and PASCAL P100 GPU with NVLink interconnection. We apply adjusting parameters in VASP INCAR file to increase performance on GPU.

2 The Basics of the Density Functional Method

We calculate energy and magnetic characteristics of Fe/Cu(111)/Fe and Co/Cu(100)/Co multilayer magnetic nanostructures. The central idea of density functional theory (DFT) [3] is to consider the electron density $n(r)$ instead of the full many-body wave functions $\Psi(r_1, \dots, r_N)$. To ensure the possibility of calculating the magnetic properties, the energy of the system is written in the form of a functional not only of the electron density $n(r)$, but also of the magnetization density $m(r)$, see formula (1). The Kohn-Sham wave functions are replaced by two-component Pauli wave functions $\Psi_\nu(r)$, capable of representing both the electron density and the magnetization density. Index ν here denotes spin states.

$$n(r) = \left\langle \Psi \left| \sum_{i=1}^N \delta(r - r_i) \right| \Psi \right\rangle, \quad m(r) = \sum_{\nu=1}^N \Psi_\nu^*(r) \sigma_\nu \Psi_\nu(r) \quad (1)$$

From the variational principle, the Kohn-Sham equations are obtained:

$$\left\{ -\frac{\hbar}{2m} \nabla^2 + V_{eff} + \sigma \cdot B_{xc}(r) - \varepsilon_\nu \right\} \Psi_\nu(r) = 0, \quad B_{xc} = \frac{\partial E_{xc}[\rho(r), m(r)]}{\partial m(r)} \quad (2)$$

Copyright © 2019 for the individual papers by the papers' authors. Copyright © 2019 for the volume as a collection by its editors. This volume and its papers are published under the Creative Commons License Attribution 4.0 International (CC BY 4.0).

In: Sergey I. Smagin, Alexander A. Zatsarinnyy (eds.): V International Conference Information Technologies and High-Performance Computing (ITHPC-2019), Khabarovsk, Russia, 16-19 Sep, 2019, published at <http://ceur-ws.org>

where B_{xc} is the effective magnetic field arising from the exchange-correlation energy.

The main problem associated with the density functional theory method is that exact analytical expressions for exchange and correlation functionals are known only for the particular case of a gas of free electrons. Nevertheless, the existing approximations allow us to calculate a number of physical quantities with sufficient accuracy.

In this work we used GGA (generalized gradient approximations) approximations in terms of Perdew–Burke–Ernzerhof (PBE) [9]:

$$E_{xc}^{GGA}[n_{\uparrow}(\vec{r}), n_{\downarrow}(\vec{r})] = \int \varepsilon[n_{\uparrow}(\vec{r}), n_{\downarrow}(\vec{r}), \nabla n_{\uparrow}(\vec{r}), \nabla n_{\downarrow}(\vec{r})] d\vec{r} \quad (3)$$

The essence of the projection augmented wave (PAW) method is to transform the pseudowave functions, obtained in the pseudopotential method into all-electron wave functions, thereby restoring the information lost when considering pseudowave functions. The number of plane wave components is limited by the Cut-off Energy. To describe the first Brillouin zone we used standard method Monkhorst–Pack with the parameter K-points characterizing regular grid in k-space [10].

3 Compiling the Parallel Version of VASP Code for Openpower and Intel Architectures

Official support of the GPU calculations appeared in VASP from version 5.4.1 and in the our work we used version 5.4.4. VASP has one precompiled configuration file, named makefile.include, with a lot of parameters. Showing all parameters is redundancy and we present the main part of it in table 1. We use Intel Parallel Studio XE C/C++ with Intel MKL library to compile the VASP package on X86_64 architecture. The optimization flags were chosen -O1 and -O2 because compilation with harder optimization was not complete successfully. It was used Ubuntu 16.04 with 4.4.0-137 kernel. We used XIC 13.1.5 and Xlf 15.1.5 with including ESSL library on CentOS 7 with 3.10.0-514 kernel to compile VASP on IBM Power System S822LC.

Table 1: Compilers options

Compiler	Version	Flags
IBM XL C/C++	13.1.5	-g -q64 -O3 -qarch=pwr8 -qtune=pwr8:st -qfullpath -qsaveopt
IBM XL Fortran	15.1.5	-g -q64 -O3 -qarch=pwr8 -qtune=pwr8:st -qfullpath -qsaveopt -qflag=i:e -qsuppress=cmpmsg
NVIDIA CUDA compilation tools	8.0.61	-DCUDA_GPU -DRPROMU_CPROJ OVERLAP -DCUFFT_MIN=28 -UscaLAPACK -fPIC -DADD -DMAGMA_WITH_MKL -DMAGMA_SETAFFINITY - DGPUSHMEM=300 -DHAVE_CUBLAS
Intel Parallel Studio XE C/C++	2017	-O2 -f_com=no -free -w0
Intel Parallel Studio XE Fortran	2017	-O1 -mkl=sequential -lstdc++

4 Hardware Information

IBM Power System S822LC is two-socket HPC system with two POWER8 CPUs with 20 cores running at 4 GHz and interconnected with two Nvidia Pascal P100 GPUs with a high bandwidth (80GByte in and 80GByte out) NVLink 1.0 interface (Fig. 1). It is very important for exchange data between multiple GPUs and fast load data from CPU. The major goal of this system is efficiently use GPU units and accelerating calculations. A large part of HPC resources installed during the last decade are based on Intel CPUs. Novel generations of Intel CPUs present a wide spectrum of multicore processors [5]. Intel Core i7 4770 is desktop processor, but it is “tock” model in Intel extensive strategy of microprocessor development – it is mostly complete 22 nm architecture. We compared IBM POWER8 with Intel Haswell because both architectures introduced in 2013 year and had 22 nm technical processes.




Figure 1: NVLink communications protocol in IBM Power System S822LC

5 Model and Simulation Parameters

In this work, the results of numerical first-principles calculations of the energy and magnetic characteristics for cobalt and iron films on a copper surface obtained by using VASP software package by means of the Projector Augmented Wave (PAW) method are presented. The values of the total energy of collinear spin configuration, the total magnetic moment and the magnetic moments of Co and Fe atoms are calculated. We investigated a system consisting of a copper slab and adsorbed on it from both sides by a ferromagnetic film with the thickness of the films in three monoatomic layers. The multilayer structure was simulated using a periodic 2×2 36-atom supercell with the lattice constant corresponding to the copper substrate $a = 3.6367$ (5) Å, which we obtained as a result of calculations taking into account the optimization of the lattice parameters. The surface face orientation is (100) for Co/Cu system and (111) for Fe/Cu system.




Figure 2: Representations of Co/Cu/Co and Fe/Cu/Fe multilayer nanostructures

For Fe/Cu system the calculations of the total energy were realized for ferromagnetic and two different antiferromagnetic spin configurations. The antiferromagnetic spin configurations for which the calculation was carried out are shown in Fig. 3. The magnetic moment of the atoms is directed along the z axis.

VASP INCAR file has several adjusting parameters that can increase GPU performance. The main of those are NCORE, NPAR, NSIM, LPLANE.

- NCORE determines how many cores work on individual orbital;
- NPAR depend on NCORE as $\text{NCORE} = \text{number of the cores} / \text{NPAR}$;

- If NPAR is equal to the number of cores than NCORE = 1, therefore, one orbital will treat by one core.
- In the INCAR file we need to set only NCORE or NPAR parameter because NPAR have precedence over NCORE and in the relatively modern version of VASP using NCORE instead of NPAR is recommended.
- NSIM is an important parameter to get calculation on GPU faster. It changes the number of bands treated simultaneously. There is an opinion that for GPU NSIM parameter needs to be increased while we have free memory on the GPU.
- LPLANE is a useful parameter for optimization which can reduce intercommunication time, but it is actual firstly for massively parallel systems, according to VASP documentation.

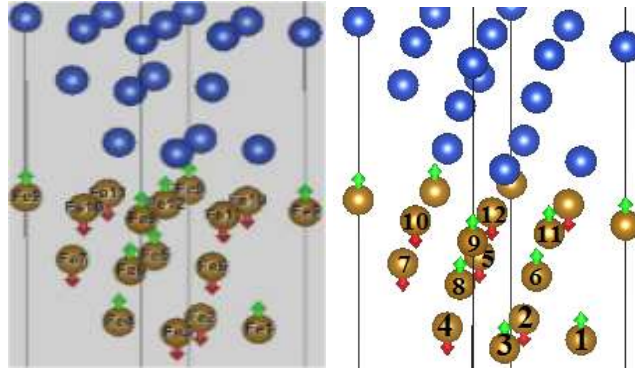


Figure 3: Fe/Cu/Fe antiferromagnetic spin configurations

Table 2: Parameters of modeling structures

	Co/Cu	Fe/Cu	Fe/Cu AF1	Fe/Cu AF2
Cutoff energy	500 eV	350 eV	350 eV	350 eV
K-points	12	10	10	10
Number of atoms	36	36	36	36
Thickness of vacuum layers	5 Å	4 Å	4 Å	4 Å

6 Estimation of Performance and Accuracy of Calculations

For CoCu system we used three configurations to compare VASP calculation times with similar INCAR parameters for IBM POWER8 CPU and Intel Haswell CPU. We set LREAL=.TRUE. as described at VASP official documentation and use NCORE=1 with one MPI thread for GPU calculations. We do not use NVIDIA MPS system and do not set NSIM parameter clearly, but we know that it is important for large tasks especially. Core i7 has only 4 real cores and we run VASP with 4 processes only. On POWER8 we run VASP on 8 cores for using most of one CPU.

Table 3: Comparison of the accuracy of the calculations performed on the CPU and GPU in VASP for CoCu

Architecture	Total magnetization	Free energy (eV)
Intel Core i7-4770 Haswell	40.646	-202.20538224
IBM POWER8	40.646	-202.20538199
IBM POWER8+Nvidia P100	40.658	-202.04101161

The results of magnetization and free energy calculations (Table 3) are well correlated. The calculations with the GPU provide less accuracy, but the value of the error is not so sufficient. The times of calculations are different for POWER8 system and Intel Core i7 (Table 4). The one POWER8 thread was more efficient than one Intel thread for the VASP calculations. If we use GPU only with one MPI thread we have much better performance (Table 4) than Intel or POWER CPUs even without optimizations of VASP parameters in INCAR file.

For FeCu systems we used NCORE = 4 and NSIM = 32 to get better performance for GPU calculations We perform simulation of the ferromagnet FeCu system on ten POWER8 cores to compare execution times with GPU

(Fig. 4). As we can see in Table 5 antiferromagnet spin configuration need sufficient more memory than a ferromagnet. Execution time on GPU highly depends on spin configuration too. It makes sense to note, that GPU utilization is not full and floating from about 20% to 70% during calculation for both systems.




Figure 4: Dependence of the execution times on used architecture and different tasks

Table 4: Execution times for different computing systems (in hours)

Architecture	Co/Cu	Fe/Cu	Fe/Cu AF1	Fe/Cu AF2
Intel Core i7-4770 Haswell	38			
IBM POWER8	14.4	9.55		
IBM POWER8+Nvidia P100	8	4.04	12.15	4.89

Table 5: Used memory for different computing systems (in Gb)

Architecture	Co/Cu	Fe/Cu	Fe/Cu AF1	Fe/Cu AF2
Intel Core i7-4770 Haswell	6.91			
IBM POWER8	6.67	4.87		
IBM POWER8+Nvidia P100	29.64	15.9	25.88	20.84

7 Conclusions

VASP is widely used by researchers to get characteristics of solids and multilayer magnetic structures. NCORE and NSIM parameters can be very useful to maximize performance on GPU. The values of acquired quantities and accuracy of GPU calculations are in a good agreement with CPU results. To use VASP efficiently with GPUs more memory and calculation time is required in comparing with calculations on CPU, especially for antiferromagnet spin configurations.

Acknowledgements

We would like to thank the IBM experts, who help us to optimize the VASP package for IBM Power Systems S822LC. This research was supported by the grants 17-02-00279, 18-42-550003 of Russian Foundation of Basic Research and by the grant MD-6868.2018.2 of the President of the Russian Federation. The simulations were supported by the computational resources of Shared Facility Center "Data Center of FEB RAS" (Khabarovsk) [10]. Computations were performed with the methods and techniques which had been developed under the RFBR scientific project number 18-29-03196.

References

1. Lejaeghere, K., Bihlmayer, G., Björkman, T., et al.: Reproducibility in density functional theory calculations of solids, *Science*. 351:aad3000 (2016)
2. Kondrashov, R.A., Mamonova, M.V., Povoroznuk, E.S, Prudnikov, V.V.: First-principles investigations of the atomic structure and magnetic properties of Ni and Co films on Cu substrate, *Lobachevskii Journal of Mathematics*. 38:940 (2017)
3. Kresse, G., Furthmüller, J.: Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set, *Phys. Rev. B*: 54:11169 (1996)
4. Kresse, G., Marsman, M., Furthmüller, J.: VASP THE GUIDE (2015) <https://cms.mpi.univie.ac.at/vasp/vasp/vasp.html>
5. Stegailov, V., Vechev, V.: Efficiency Analysis of Intel and AMD x86 64 Architectures for Ab Initio Calculations: A Case Study of VASP, In: Voevodin, V., Sobolev, S.: (eds) *Supercomputing RuSCDays 2017. Communications in Computer and Information Science*, vol 793. Springer, Cham (2017)
6. Giannozzi, P., Baroni, S., Bonini, N., et al.: QUANTUM ESPRESSO: a modular and open-source software project for quantum simulations of materials, *Journal of Physics: Condensed Matter*. 21:395502 (2009)
7. Gonze, X., Amadon, B., Anglade, P.M., et al.: ABINIT: First-principles approach to material and nanosystem properties, *Comput. Phys. Commun.* 180:2582 (2009)
8. Schwarz K., Blaha P.: Solid state calculations using WIEN2k, *Computational Materials Science*. 28: 259–273 (2003)
8. Perdew, J.P., Burke, K., Ernzerhof, M.: Generalized Gradient Approximation Made Simple, *Phys. Rev. Lett.* 77:3865 (1996)
9. Monkhorst, H.J., Pack, J.D.: Special points for Brillouin-zone integrations, *Phys. Rev. B* 13:5188 (1976)
10. Sorokin, A.A., Makogonov, S.I., Korolev, S.P.: *Scientific and Technical Information Processing* 4:302 (2017)