Computer Physics Communications 196 (2015) 267–275

Contents lists available at ScienceDirect

Computer Physics Communications

journal homepage: www.elsevier.com/locate/cpc

Automated generation of highly accurate, efficient and transferable pseudopotentials

R.A. Hansel [a], C.N. Brock [a], B.C. Paikoff [b], A.R. Tackett [c], D.G. Walker [b,*]

[a] Materials Science Program, Vanderbilt University, United States
[b] Department of Mechanical Engineering, Vanderbilt University, United States
[c] Department of Astronomy and Physics, Vanderbilt University, United States

ARTICLE INFO

Article history:
Received 31 December 2014
Received in revised form
15 June 2015
Accepted 17 June 2015
Available online 26 June 2015

Keywords:
DFT
Pseudopotential
Optimization

ABSTRACT

A multi-objective genetic algorithm (MOGA) was used to automate a search for optimized pseudopotential parameters. Pseudopotentials were generated using the `atomPAW` program and density functional theory (DFT) simulations were conducted using the `pwPAW` program. The optimized parameters were the cutoff radius and projector energies for the *s* and *p* orbitals. The two objectives were low pseudopotential error and low computational work requirements. The error was determined from (1) the root mean square difference between the all-electron and pseudized-electron log derivative, (2) the calculated lattice constant versus reference data of Holzwarth et al., and (3) the calculated bulk modulus versus reference potentials. The computational work was defined as the number of flops required to perform the DFT simulation. Pseudopotential transferability was encouraged by optimizing each element in different lattices: (1) nitrogen in GaN, AlN, and YN, (2) oxygen in NO, ZnO, and SiO_4, and (3) fluorine in LiF, NaF, and KF. The optimal solutions were equivalent in error and required significantly less computational work than the reference data. This proof-of-concept study demonstrates that the combination of MOGA and *ab-initio* simulations is a powerful tool that can generate a set of transferable potentials with a trade-off between accuracy (error) and computational efficiency (work).

1. Introduction

Pseudopotentials with density functional theory have long been regarded as an efficient approach to perform electronic structure calculations compared to all-electron calculations [1] because (1) the number of orbitals to be calculated is reduced, and (2) the number of basis functions required to describe a smooth potential, compared to the all-electron counterpart, is reduced. In addition, relativistic effects can be incorporated into the pseudopotential so non-relativistic calculations can reproduce relativistic effects [2]. Due to these benefits, researchers have been searching for pseudopotentials representing each element in the periodic table that are accurate, efficient, and transferable. Accuracy refers to how well the pseudized solution matches the all-electron solution. Efficient refers to how fast the calculation proceeds, which is often expressed in terms of the plane-wave cut-off. Transferability refers to the universality of the pseudopotential in multiple lattices or molecules [3].

With efficiency in mind, pseudopotentials were first used by Phillips and Kleinman [4] as early as 1959. Perhaps a more complete treatment was provided by Hamann in 1979 [5] where the density norm was equated between the pseudized density and the all-electron density. In 1982, Bachelet et al. [6] recount the seminal accomplishments that lead to their library of coefficients for norm-conserving pseudopotentials of all elements up to plutonium. The transferability was promoted by matching log derivatives, but little effort was devoted to verifying the transferability in other systems, and accuracy was assumed since the pseudopotential was generated from the all-electron solution. In addition, by virtue of the reduction in orbitals, the pseudized solution was known to be vastly superior to the all-electron solution in terms of computational speed and scalability, so no attention was devoted to tuning the potentials for efficiency. Unfortunately, these potentials require a large plane-wave basis set and, consequently, are not very efficient.

To improve efficiency, Vanderbilt relaxed the norm-conserving constraint by adding charge to the core to match the all-electron density solution [7]. This new approach (resulting in "ultra-soft" potentials) indeed improved the efficiency, but at the expense of additional parameters to insure charge conservation. Shortly

* Corresponding author.
E-mail address: greg.walker@vanderbilt.edu (D.G. Walker).

http://dx.doi.org/10.1016/j.cpc.2015.06.014
0010-4655/

afterward, Blöchl introduced the PAW pseudopotential based on the all-electron wave-function (as opposed to the electron density) [8]. This approach has the distinct advantage that electron wave functions inside the cut-off radius can be recovered and that the size of the set of parameters required to describe the potential are of the order of the norm-conserving potentials. Despite the apparent difference in the approaches, the pseudopotentials created from each method are remarkably similar, and Kresse showed how they are mathematically related [9]. In fact, the accuracy, efficiency, and transferability are also comparable [10]. Still few efforts were made to improve efficiency, to increase accuracy, and to ensure transferability simultaneously.

Today materials exploration is being performed by scaling up DFT calculations to thousands of DFT test cases to generate libraries of materials properties from DFT analysis [11]. With such large numbers of tests, researchers are finding that results are sensitive to pseudopotential selection, and conservative choices in pseudopotential creation are often made to insure transferability and accuracy [12]. The PAW functionals that are being proposed are with reduced number of core electrons to improve transferability, which increases the computation time of the DFT, not because of increased plane-wave basis sets, but because additional orbitals must be computed. Regardless of the reason, the computation time of the calculations increases. Interestingly, despite the obvious drawbacks of norm-conserving pseudopotentials, some researchers have been using them for high-throughput DFT simulations because of their inherent transferability [13] without significant reduction in efficiency. Nevertheless, efficiency remains an important issue in pseudopotential design.

In general, quality pseudopotentials are developed by hand by experts and placed in a library for subsequent use by users of DFT tools. Recently several libraries of pseudopotentials have been developed. Garrity et al. [10] have developed a library using the Vanderbilt pseudopotential code and compared it to VASP [9], which is proprietary, and PSLIB [14], which is open-source but still under development. The `abinit` tool developers maintain a library of pseudopotentials built by individuals and posted for all to evaluate. Holzwarth also maintains a library generated by the `atomPAW` code [15,16]. None of these libraries claims to provide accurate pseudopotentials that are also efficient and transferable.

This ad-hoc approach, as described in Ref. [10], leads to pseudopotentials that are taken as standard when their applicability to new systems and alternative structures (transferability) can be dubious. For example, the maintainers of the `abinit` tool and the `atomPAW` tool hold repositories of pseudopotentials for many elements in the periodic table [17,18]. Some potentials were developed to be transferable to other crystal systems and have high accuracy at the expense of efficiency. On the other hand, pseudopotentials with high efficiency can be designed for use in large crystal systems; however, the accuracy of these potentials is questionable. Actually, these potentials are expected to work for the scant systems that were tested or for essentially similar crystal systems [19], and so users are urged to check the potentials for themselves.

Physical parameters such as the electron configuration (how many electrons to pseudize) and exchange–correlation method can further confound pseudopotential development because pseudopotentials developed for one configuration may not work for another. In addition, the mathematical or simulation parameters that define the pseudopotential, such as cut-off radius and projector energy, can significantly affect the accuracy of a pseudopotential across various crystal systems. Consequently, generating accurate, transferable, and computationally efficient pseudopotentials is a daunting task, which requires extensive hand-tuning and copious computational resources.

The goal of this study is to demonstrate that accurate, computationally efficient, and transferable pseudopotentials can be generated using global optimization techniques. In this work, we have confronted these challenges by using a multi-objective genetic algorithm (MOGA) to automate the search for optimized pseudopotential parameters.

Genetic algorithms (GA) are adept at global optimization of complex engineering problems with coupled parameters and multiple-minima design spaces [20]. These algorithms can be highly parallelized and are extremely powerful due to the ability to automate the exploration of extremely large search spaces. Our primary interest in this approach, however, is because the GA is particularly suited to multi-objective optimization [21,22]. This proof-of-concept study shows that combining the GA with `ab-initio` simulations can produce a Pareto set of solutions that show the trade-off between accuracy and computational efficiency. The optimal solutions generated here are equivalent in accuracy and are significantly more computationally efficient compared to reference data of Holzwarth et al. [23]. Also, the pseudopotentials were tested in three different crystal systems to establish the transferability of the potential. Ultimately, the work presented here shows that using a heuristic approach, such as MOGA, to generate viable pseudopotentials produces valuable results and can make a significant contribution to the study of electronic structure.

Furthermore, this approach could be expanded to generate a database of pseudopotentials for the entire periodic table of elements with pre-determined objectives. Ideally, these potentials would be accurate, highly transferable, and computationally efficient. The technique is equally applicable to other types of pseudopotentials such as norm-conserving [5] or ultra-soft [7].

2. Method

The multi-objective genetic algorithm (MOGA) from the Design Analysis Kit for Optimization and Terascale Applications (DAKOTA) software package was used in this work. DAKOTA was developed at Sandia National Laboratory, and a description of the MOGA technique can be found in the DAKOTA Reference Manual [24].

Fig. 1 shows a schematic diagram of the genetic algorithm structure and how the pseudopotential and DFT tools are integrated into the optimization. Although a complete description of genetic algorithms can be found in Ref. [25], basic concepts as they relate to electronic structure calculation are provided here. The algorithm begins by generating a set of input parameters for each pseudopotential to be optimized. In our case, the parameters that are varied include the cut-off radius and projector energies pertinent to the particular atom. Each set of parameters that constitute the potential is called an individual or design point, and many design points, randomly distributed within a specified range of values, are created. The set of design points is called a generation. So a generation is really a set of different pseudopotentials that could be used in a DFT calculation. Next, the quality of each individual is evaluated based on two objectives: (1) the computational work required to complete a DFT simulation and (2) the error of the pseudopotential. The error objective is defined in Sections 2.1 and 2.2, but both objectives are found by running a DFT analysis using each design point. Once all individuals have been evaluated, a portion of the design points that represent the highest quality points are combined along with additional randomness to create an entirely unique set (or generation) of design points to be tested.

The algorithm terminates when either a specified number of generations (or design points) has been produced, or when design points have reached a satisfactory quality or fitness level. For this study, the fitness convergence criterion for the non-dominated (highest quality) solutions was based on three metrics. The first metric evaluates the expanse of the non-dominated individuals. The expansion metric is computed by tracking the design points with the non-dominated solutions. Any movement of the highest quality design points is noticed and the maximum percentage