



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Large-scale quantum mechanical simulations of high-Z metals

L. H. Yang, R. Hood, J. Pask, J. Klepeis

January 19, 2007

Journal of Computer-Aided Materials Design

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Large-scale quantum mechanical simulations of high- Z metals

L. H. Yang,* R. Q. Hood, J. E. Pask, and J. E. Klepeis

*Lawrence Livermore National Laboratory,
University of California, Livermore, California 94550*

(Dated: January 16, 2007)

Abstract

High- Z metals constitute a particular challenge for large-scale *ab initio* calculations, as they require high resolution due to the presence of strongly localized states and require many eigenstates to be computed due to the large number of electrons and need to accurately resolve the Fermi surface. Here, we report recent findings on high- Z materials, using an efficient massively parallel planewave implementation on some of the largest computational architectures currently available. We discuss the particular architectures employed and methodological advances required to harness them effectively. We present a pair-correlation function for U, calculated using quantum molecular dynamics, and discuss relaxations of Pu atoms in the vicinity of defects in aged and alloyed Pu. We find that the self-irradiation associated with aging has a negligible effect on the compressibility of Pu relative to other factors such as alloying.

*Electronic address: yang1@llnl.gov

I. INTRODUCTION

Over the course of the past few decades, density functional theory (DFT) [1, 2] has proven to be an accurate and reliable foundation for the understanding and prediction of a wide range of materials properties, from the first principles of quantum mechanics (*ab initio*), with no empirical or adjustable parameters. The parameter free, quantum mechanical nature of the theory facilitates both fundamental understanding and robust predictions across the gamut of materials systems, from insulating organics to metallic actinides.

However, the solution of the equations of DFT is a formidable task and this has severely limited the range of materials systems which can be investigated by such rigorous, quantum mechanical means. High- Z metals (i.e., those of high atomic number) constitute a particular challenge for large-scale *ab initio* calculations, as they require high resolution due to the presence of strongly localized states and require many eigenstates to be computed due to the large number of electrons and need to accurately resolve the Fermi surface.

Of paramount importance in the investigation of new and unknown materials, and of more common materials under extreme conditions, is the generality and accuracy of the theoretical methods employed. First and foremost, a quantum mechanical description is required whenever departures from isolated-atomic or known condensed-matter configurations may be significant. However, a merely quantum mechanical approach is not enough in the investigation of such unknown systems: the approach must be general, equally applicable to all atomic species and configurations, and systematically improvable so that errors can be clearly known and strictly controlled.

The planewave (PW) pseudopotential method [3] is among the most widely used *ab initio* methods which afford this level of generality and systematic improvability. The accuracy and generality of the PW method arises from its nature and basis: a variational expansion approach, in which solutions are represented as a linear combination of basis functions, employing a Fourier basis. Due to the completeness of the basis, any condensed matter system can be modeled with arbitrary accuracy, in principle, by simply adding sufficient wavenumbers to the basis. In addition, the PW basis has a number of desirable properties which facilitate computations, such as orthonormality and independence of atomic positions. However, the global, uniform nature of the PW basis leads to substantial inefficiencies in large-scale calculations involving localized states, such as those considered here involving high- Z metals. Because the resolution of the basis is strictly uniform at all points, the need for higher resolution in the vicinity of the nuclei requires more basis functions at all points, leading to increased computation. And because every basis function overlaps every other, computations are difficult to localize, leading to difficulties in parallel implementation.

These difficulties have inspired extensive research in the past decade into new “real space” approaches, such as finite-difference and finite-element based approaches [4–6], which solve the required quantum mechanical equations in a strictly local representation, and so allow both variable resolution and natural parallelization. They have also inspired much work on the key bottlenecks in the standard PW approach, leading to significant progress on these fronts. Recent progress on the problem of resolution includes ultrasoft pseudopotentials [7], optimized pseudopotentials [8, 9], adaptive coordinate transformations [10–12], and the projector augmented wave method [13]. Recent progress on parallelization includes new fast Fourier transform formulations [14, 15], localized PW formulations [16], and efficient distribution schemes for electronic states and PW coefficients on large numbers of processors [17].

In this paper, we discuss our large-scale PW pseudopotential simulations of high- Z metals.

In Sec. II, we discuss the computational platforms employed. In Sec. III, we discuss the algorithmic advances we have implemented to enable such simulations on these platforms. In Sec. IV, we discuss results for U and Pu, the largest of these calculations employing 16,384 processors on the massively parallel BlueGene/L platform. We summarize in Sec. V.

II. COMPUTATIONAL PLATFORMS

Pu calculations were performed on the BlueGene/L (BG/L) platform [18] at LLNL. BG/L uses a large number of low power processors coupled with powerful interconnects and communication schemes. It consists of 65,536 compute nodes, each with two IBM PowerPC 440 processors (at 700 MHz clock speeds) and 512 MB of shared memory. The theoretical peak performance is 5.6 Gflops per node, or 367 Tflops for the full machine. Each processor has 32 KB instruction and data caches, a 2 KB L2 cache, and a 4 MB L3 cache, which is shared with the other processor on the node. Each node has two floating-point units that can perform fused multiply-add operations. In its default mode (coprocessor mode), one of the processors on the node manages the computation while the other manages communication. In an alternate mode (virtual mode), both processors can be used for computation. The system uses a highly optimized lightweight Linux distribution and does not allow access to individual nodes. The nodes are interconnected through multiple complementary high-speed low-latency networks, including a 3D torus network and tree network. The nodes are connected as a $64 \times 32 \times 32$ 3D torus, which is used for common interprocessor communications. The tree network is optimized for collective operations such as broadcast and gather. The point-to-point bandwidth of the 3D torus network is 175 MB/s per link and 350 MB/s for the tree network.

U calculations were performed on the ASC Q platform [19] at LANL. The QA and QB segments consist of 1,024 HP/Compaq AlphaServer ES45 nodes, each with 4 EV68 1.25 GHz, 16MB cache processors: 768 nodes with 8 GB memory, 192 with 16 GB, and 64 with 32 GB, with Quadrics Elan3 interconnect. Theoretical peak performance for each segment of 4,096 processors is 10.2 TFlops.

III. COMPUTATIONAL METHODS

Our electronic structure method is based on the formulation of the nonlinear Kohn-Sham (KS) equations [2] as a minimization problem for a total energy functional $E_{\text{total}}[\Psi_s, \tau_I, \alpha_v]$, where Ψ_s represents the electronic degrees of freedom (the wavefunctions), τ_I represents the ionic coordinates, and α_v represents possible constraints imposed on the system (volume, strain, etc.). The minimization of the functional uses gradient vectors and corresponding conjugate components. Our parallel algorithm for the electronic degrees of freedom is based on the band-by-band serial algorithm proposed by Teter, Payne, and Allan [20], modified to facilitate the treatment of large unit cells and metallic systems. The PPMD (Petascale Pseudopotential Molecular Dynamics) code allows the optimization of electronic and structural properties at zero and finite temperatures. The parallel algorithms employed in the current PPMD code are an extension of those employed in the previous PCGPP code [21]. In this section, we discuss the main features of the current PPMD implementation which enable the effective use of massively parallel computational platforms. A more detailed description of the code structure and algorithms will appear elsewhere [22].




FIG. 1: Hybrid decomposition scheme employed in PPMD code. Wavefunctions $\Psi_s(\mathbf{r})$ are discretized in real- and reciprocal-space and decomposed by \mathbf{k} -point (\mathbf{k}), band (b), and domain (p) in such a way as to maximize data locality for a given node geometry.

For fixed ionic coordinates and lattice vectors, the planewave DFT problem is formulated as the minimization of the energy functional $E_{\text{total}}[\Psi_s]$ with respect to electronic wavefunctions (or KS orbitals) $\Psi_s, s = 1 \dots n$, subject to orthonormality constraints. According to band theory, for a periodic solid, the index s actually contains subindices, \mathbf{k} and n , where n is the band index and \mathbf{k} is the crystal momentum. In the planewave basis, the wavefunction is expressed as $\Psi_s(\mathbf{r}) = \exp(i\mathbf{k} \cdot \mathbf{r}) \sum_{\mathbf{G}} C_{\mathbf{G}} \exp(i\mathbf{G} \cdot \mathbf{r})$, where \mathbf{G} are the reciprocal lattice vectors and $C_{\mathbf{G}}$ are the Fourier coefficients. A typical large-scale problem for high- Z metals will have a few tens of vectors \mathbf{k} , on the order of thousands of bands n , and over a million coefficients $C_{\mathbf{G}}$ for each n and \mathbf{k} . For the data structure of the wavefunctions, a number of data distribution strategies for massively parallel algorithms can be considered. Our focus on ever larger and more complex high- Z metallic systems requires that the implementation be *flexible* and *portable*, without sacrificing significant efficiency, to facilitate full and immediate exploitation of current and next-generation large-scale platforms as they become available. Therefore, the decomposition of the data structure has to be general and suitable for these purposes.

A. General features

The current PPMD code can use any one or a combination (hybrid) of three decompositions. A schematic representation is shown in Fig. 1.

1. **\mathbf{k} -point decomposition (“k-decomposition”)**: The bulk of computations performed at each \mathbf{k} -point are independent. This could be exploited with a relatively small number of nodes with large memory per node (8 GB or more). Because traditional electronic-structure algorithms have employed a serial loop over \mathbf{k} -points, each eigensystem solution being independent of the others, this is a much simpler approach and better exploits data locality than the other approaches below. However, the largest machines are not of this type. Rather, like BG/L, they tend to have many nodes with relatively small memory per node. In addition, even with an appropriate machine, this approach can exhibit stability problems, requiring great care in SCF mixing to attain convergence [20, 21]. The approach used in PPMD is based on the band-by-band algorithm of Ref. [20], with modifications to improve convergence for metallic systems and increase efficiency in parallel implementation. For a k-decomposition, where all

bands for a subset of \mathbf{k} s are contained on the same node, all bands from one \mathbf{k} on each node are updated and summed to update the density and potential, then all bands from the next \mathbf{k} on each node are incorporated, then all bands from the next, etc., until all bands from all \mathbf{k} s are incorporated.

2. Band decomposition (“b-decomposition”): The computations performed with respect to different electronic bands at the same \mathbf{k} -point are largely independent and can therefore be performed simultaneously on different nodes. This could be exploited by allocating enough nodes to cover a small number of bands per node. However, this would pose problems of both flexibility and communication expense because allocating the exact number of nodes for the problem size may not be possible and communication cost increases as the number of bands per node decreases.
3. Domain decomposition (“p-decomposition”): This is the natural way to parallelize a planewave code if a highly efficient parallel 3D FFT is available. In this scheme, the \mathbf{r} (real space) and \mathbf{G} (Fourier space) components of the wavefunctions and charge density are distributed evenly over the entire machine. This is the most flexible approach with less restriction on the number of nodes needed for the wavefunction distribution. In addition, with this scheme, a modified band-by-band algorithm can be used in the minimization of the energy functional. In PPMD, the full charge density is updated by a mixture of the current and previous partial densities $|\Psi_{n,k}|^2$. However, as the number of nodes across which the wavefunctions and density are distributed increases, the cost of communication increases. For this reason, we restrict the distribution to 16 nodes or less in practice.

The PPMD code uses a combination of the above k, b, and p decompositions to optimize data locality and node assignment on a given platform. A search algorithm [22] to automatically determine efficient node assignments has been implemented. The PPMD parallel implementation involves four key components: 3D complex FFT, parallel I/O, symmetrization of the charge density, and wavefunction distribution. Here we discuss the two most critical components: 3D complex FFT and wavefunction distribution.

1. Distributed 3D complex fast Fourier transforms: In the PPMD code, FFTs are employed as follows:
 - Perform 3D FFT on array X from \mathbf{G} -space to \mathbf{r} -space.
 - Do calculations on array X in \mathbf{r} -space.
 - Perform 3D FFT on array X from \mathbf{r} -space to \mathbf{G} -space.

The 3D FFT transforms the array X from \mathbf{G} -space to \mathbf{r} -space then back to \mathbf{G} -space after some calculations have been performed in \mathbf{r} -space, where they are better localized. There is no need to keep the order of array indices in \mathbf{r} -space the same as in \mathbf{G} -space. As a matter of fact, in this case we can express a 3D FFT as three sets of 1D transforms corresponding to x , y , and z directions. Each of the 1D transforms in a set is independent of the others in that set and so can be performed in parallel. The parallel FFT routines in PPMD require less memory, and are thus more suitable for machines such as BG/L, than those found in the standard parallel libraries such as FFTW [23]. They also allow user control of node assignments. Table I illustrates the

TABLE I: Three steps to transform a 3D complex array from \mathbf{r} -space to \mathbf{G} -space. The total number of nodes $N = N_1 \times N_2$. Each step shows the distribution of the data on the nodes resulting from each operation.

| Step | x | y | z | Operations |
|--------|-------|-------|-------|--|
| 1 | N_1 | N_2 | 1 | 1D FFT along z |
| 2 | N_1 | 1 | N_2 | Transpose y - z , 1D FFT along y |
| 3 | 1 | N_1 | N_2 | Transpose x - y , 1D FFT along x |
| Output | 1 | N_1 | N_2 | \mathbf{G} -space order: z, y, x |

TABLE II: Timing results for 128-atom U supercell on BG/L. Shown is the number of nodes, time in FFT operations, time in communications, and total time to complete a full self-consistent calculation for each parallel decomposition. All times are in seconds.

| Parallel decomposition | Number of nodes | FFT time | Communication time | Total time |
|------------------------|-----------------|----------|--------------------|------------|
| (k,b) | 4,096 | 10,556 | 8,394 | 25,670 |
| (k,b) | 8,192 | 5,820 | 4,215 | 13,254 |
| (k,p) | 8,192 | 4,215 | 4,344 | 18,593 |
| (b,p) | 8,192 | 3,820 | 4,200 | 14,859 |
| (k,b,p) | 16,384 | 3,843 | 4,680 | 8,211 |

data flow of the reduced FFT operations for a transformation from \mathbf{r} -space to \mathbf{G} -space. The data arrangement in \mathbf{r} -space is local along the z -direction but distributed across the x - y planes, while the arrangement in \mathbf{G} -space is local along the x -direction but distributed across the y - z planes. Only two transpose operations and three 1D FFTs are needed to transform the full 3D data from \mathbf{r} -space to \mathbf{G} -space.

2. Data layout for wavefunctions: If the wavefunction coefficients $C_{\mathbf{G}}$ are distributed across the nodes, the two transpose operations described above require global communications and so can become bottlenecks. Our communication scheme in this case only requires non-zero data elements to be exchanged between nodes. In order to minimize the amount of communications and thus maximize the data locality, we also restrict each wavefunction to be distributed on no more than 16 nodes.

B. Code performance

Table II shows timing results for PPMD for a 128-atom bulk U supercell on BG/L. Each U has 14 electrons in valence. A total of 2,048 bands and 4 reduced \mathbf{k} -points in the Brillouin zone were included to calculate the SCF charge density. A PW cutoff of 100 Ry and FFT mesh of $128 \times 128 \times 128$ points were used. Each benchmark run represents a complete self-consistent calculation of about 30 PCG outer-loop steps, with 3 inner-loop steps per PCG outer-loop step. Four hybrid decomposition schemes were tested: (k,b), (k,p), (b,p), and (k,b,p). Up to 16,384 nodes were used. For this case, we found the most efficient approach to be (k,b) decomposition, as expected due to its better data locality and correspondingly lower communications overhead. Although there are two processors per node on BG/L, we

TABLE III: Ground-state properties of α -U calculated using a planewave pseudopotential approach and FP-LTMO (Ref. 27). The experimental results (Ref. 28) were measured at room temperature. The equilibrium volume and the lattice constants a, b, and c are in units of a.u. The bulk modulus is in GPa.

| | pseudopotential | FP-LMTO | Experiment |
|--------|-----------------|---------|----------------------------|
| volume | 139.8 | 139.5 | 140.16 |
| a | 5.45 | 5.376 | 5.3957 |
| b | 10.8 | 10.99 | 11.093 |
| c | 9.53 | 9.441 | 9.3670 |
| y | 0.10 | 0.1025 | 0.102 |
| B | 133.5 | 133.0 | 104(2), 135.5 ^a |
| B' | 5.0 | 5.4 | 6.2 |

^aRef. 29

only utilized one per node due to the limited memory (0.5 GB) on each.

IV. RESULTS AND DISCUSSION

A. Uranium quantum molecular dynamics

Quantum molecular dynamics (QMD) based upon DFT and making use of planewave *ab initio* pseudopotentials has been very successful in describing the properties of a number of different materials including simple metals and transition metals that have close-packed high-symmetry structures, such as body-centered cubic (bcc), face-centered cubic (fcc), and hexagonal close-packed (hcp). It has to the best of our knowledge not yet been used to study metals in which *f* electrons participate in the bonding as in the light actinide metals, Th through Pu. These 5*f* electrons have a dominant influence upon the electronic structure imparting highly directional bonding and producing polymorphism and low symmetry crystal structures not seen by other metals in the periodic table [24–26].

We constructed a planewave based pseudopotential for uranium within the DFT [1, 2] approach using a parameter independent generalized gradient approximation (GGA) [30] for the exchange-correlation term by solving for the all-electron uranium atom U^{2+} in the reference state $6s^26p^66d^15f^37s^0$. The pseudopotential had a norm-conserving nonlocal Troullier-Martins form [31] with a total of 14 valence electrons generated using a PW cutoff of 100 Ry. Table III lists our calculated ground-state properties of α -U compared with full-potential linear muffin tin orbital (FP-LMTO) calculations [27] and experimental data. Our calculated lattice constants differ by at most 3% in comparison with experimental values.

We used Born-Oppenheimer QMD in which the low lying single-particle electronic eigenstates were computed by solving the self-consistent Kohn-Sham equations. We performed NVT ensemble QMD using a Nose-Hoover thermostat [32, 33] with 54 uranium atoms in a cubic simulation cell with periodic boundary conditions, an atomic volume of 118.3 a.u.³, a temperature of 1600 K and a time-step of 1.2 fs. The starting ionic configuration was generated by performing classical molecular dynamics simulations using a potential based on generalized pseudopotential theory, that was able to reproduce [34] the experimental




FIG. 2: Radial distribution function $g(r)$ for uranium calculated using QMD with an atomic volume of 138 a.u.^3 at a temperature of $1600 \pm 200 \text{ K}$ and pressure $4.8 \pm 0.4 \text{ GPa}$.

melting curve. The configuration was thermalized for 0.5 ps and statistics were collected for 0.5 ps. In Fig. 2, we show the calculated radial distribution function $g(r)$ for uranium. This information along with configurations taken from the QMD simulation can aid in the development of accurate interatomic potentials [35] that are valid at high temperatures and pressures, and can be used to study mechanical properties of uranium.

B. Aged Pu and Pu-Ga alloy

At ambient conditions, Pu exists in its α phase, a complex monoclinic structure with 16-atom unit cell [36]. Because Pu lies at the boundary between localized and delocalized f -electron states (between Am and Np, respectively), its electronic structure is notoriously sensitive to small perturbations and calculations can be correspondingly difficult to converge. Moreover, it is a metal with a large number of valence electrons (16 distributed among s , p , d , and f states, omitting semicore $5d$ states; 26, including semicore $5d$'s). Due to the low-symmetry, sensitive electronic-structure, and large number of electrons in valence, some highly localized, Pu presents a particular challenge for large-scale quantum mechanical calculations. In our calculations, we employed norm-conserving Troullier-Martins pseudopotentials [31] with 14 and 16 electrons in valence for U and Pu, respectively. This necessitated cutoffs of up to 120 Ry to attain sufficient convergence. Exchange and correlation interactions were modelled in the Perdew-Burke-Ernzerhof generalized gradient approximation [30].

1. Aged Pu

In order to study the effects of self-irradiation associated with aging, we modelled aged α -Pu by a doubled primitive cell with interstitial He at the most open position, as determined by Voronoi analysis, and substitutional U at the site farthest from the He to minimize defect-defect interactions (Fig. 3). Upon relaxation, we found that the interstitial He displaces from its initial position at the Voronoi vertex and its neighboring Pu atoms displace outward. The




FIG. 3: Aged α -Pu unit cell. Pu (green), interstitial He (yellow), and substitutional U (gray).




FIG. 4: Relaxed aged α -Pu cell. Pu atoms (green) relax outward, away from both interstitial He (yellow) and substitutional U (gray). White atoms represent initial (unrelaxed) positions.

four nearest neighbor Pu's move from an initial distance of ≈ 3.8 a.u. to a range of distances from $\approx 3.9 - 4.8$ a.u. The substitutional U also displaces from its initial position and its neighboring Pu's also displace outward, indicating weaker U-Pu than Pu-Pu bonds at that site. The four nearest Pu atoms move from $\approx 4.8 - 5.0$ a.u. away to $\approx 6.2 - 6.3$ a.u. upon relaxation. The resulting structure is shown in Fig. 4.

The above concentration of U-He defects corresponds to an age of approximately 760 years. In order to reach younger ages, lower concentrations must be modelled, requiring larger unit cells. To study the effects of age, a series of models were constructed containing up to 129 atoms and fully relaxed at a series of volumes to determine compressibility. To accomplish the largest of these calculations, the massively parallel PPMD code was employed on up to 8,192 processors of the BG/L machine at LLNL. Consistency of the PW results with

all-electron LMTO [37] calculations was verified for the smallest models without quantum mechanical relaxation. The calculations revealed that the self-irradiation associated with aging has a negligible effect on the compressibility of Pu relative to other factors such as alloying.

2. Pu-Ga alloy

In order to study the effects of alloying with Ga, we modelled α -Pu-Ga alloy by two doubled α -Pu primitive cells: one with substitutional Ga at the least open, highest-energy site-I, and the other with substitutional Ga at the most open, lowest-energy site-VIII. These sites were chosen to bracket the effects in the random alloy with a minimum of computation. Upon relaxation, we found, as expected, that the nearest neighbor Pu atoms relax outward, away from the Ga at the least open, highest-energy site I. More interestingly, however, we found that the neighboring Pu atoms also relax outward away from the Ga at the most open, lowest-energy site VIII. This is particularly interesting in light of experiments [38] and corresponding *ab initio* calculations [39] indicating that the partial molar volume of Ga at site VIII is comparable to that of Pu (about 20 Å³). It is not necessarily inconsistent with these results, however, as the outward relaxations may be compensated by a decrease in the local atomic volume of more distant neighbors, as also found in Ref. [39].

V. SUMMARY

We have computed a variety of properties of U and Pu liquids and solids using the massively parallel PPMD quantum molecular dynamics code on the LANL Q and LLNL BG/L machines, some of the largest computational platforms available today. We discussed the hybrid k, b, and p (**k**-point, band, and wavefunction) parallel data distribution scheme employed in the PPMD code in order to maximize data locality for a given node geometry, and demonstrated its effectiveness in tests on up to 16,384 processors. We presented a pair-correlation function for U, calculated using quantum molecular dynamics, and discussed relaxations of Pu atoms in the vicinity of defects in aged and alloyed Pu, finding that the Pu atoms relax outward away from all defects studied: He interstitials, U substitutions, and Ga substitutions alike. Finally, we found that the self-irradiation associated with aging has a negligible effect on the compressibility of Pu relative to other factors such as alloying with Ga.

Acknowledgments

Figures 3 and 4 were produced using the open-source XCrySDen visualization package [40], which is gratefully acknowledged. This work was performed under the auspices of the U.S. Department of Energy by University of California, Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

[1] P. Hohenberg and W. Kohn, Phys. Rev. **136**, B864 (1964).

- [2] W. Kohn and L. J. Sham, *Phys. Rev.* **140**, A1133 (1965).
- [3] W. E. Pickett, *Computer Physics Reports* **9**, 115 (1989).
- [4] T. L. Beck, *Rev. Mod. Phys.* **72**, 1041 (2000).
- [5] T. Torsti, T. Eirola, J. Enkovaara, T. Hakala, P. Havu, V. Havu, T. Höynälänmaa, J. Ignatius, M. Lyly, I. Makkonen, et al., *phys. stat. sol. (b)* **243**, 1016 (2006).
- [6] J. E. Pask and P. A. Sterne, *Model. Simul. Mater. Sci. Eng.* **13**, R71 (2005).
- [7] D. Vanderbilt, *Phys. Rev. B* **41**, 7892 (1990).
- [8] A. M. Rappe, K. M. Rabe, E. Kaxiras, and J. D. Joannopoulos, *Phys. Rev. B* **41**, 1227 (1990).
- [9] J. S. Lin, A. Qteish, M. C. Payne, and V. Heine, *Phys. Rev. B* **47**, 4174 (1993).
- [10] F. Gygi, *Europhys. Lett.* **19**, 617 (1992).
- [11] A. Devenyi, K. Cho, T. A. Arias, and J. D. Joannopoulos, *Phys. Rev. B* **49**, 13373 (1994).
- [12] D. R. Hamann, *Phys. Rev. B* **51**, 9508 (1995).
- [13] P. E. Blöchl, *Phys. Rev. B* **50**, 17953 (1994).
- [14] S. Goedecker, M. Boulet, and T. Deutsch, *Comput. Phys. Commun.* **154**, 105 (2003).
- [15] A. Canning and D. Raczkowski, *Comput. Phys. Commun.* **169**, 449 (2005).
- [16] C. K. Skylaris, P. D. Haynes, A. A. Mostofi, and M. C. Payne, *J. Chem. Phys.* **122**, 084119 (2005).
- [17] F. Gygi, E. W. Draeger, M. Schulz, B. R. de Supinski, J. A. Gunnels, V. Austel, J. C. Sexton, F. Franchetti, S. Kral, C. W. Ueberhuber, et al., In *Proceedings of Supercomputing 2006* (2006), international Conference on High Performance Computing, Network, Storage, and Analysis. 2006 Gordon Bell Prize winner (Peak Performance).
- [18] See, e.g., <http://www.llnl.gov/ASC/platforms/bluegenel>.
- [19] See, e.g., http://www.llnl.gov/asc/computing_resources/lanl.q.
- [20] M. P. Teter, M. Payne, and D. Allan, *Physical Review B* **40**, 12255 (1989).
- [21] L. H. Yang, in *Industrial Strength Parallel Computing*, edited by A. E. Koniges (Morgan Kaufmann, 2000), vol. X, p. 297.
- [22] L. H. Yang (2007), unpublished.
- [23] See, e.g., <http://www.fft.org/>.
- [24] J. J. Katz, L. R. Morss, and G. T. Seaborg, *The Chemistry of the Actinide Elements* (Chapman and Hall, 1986), pp. 1121–1195, 2nd ed.
- [25] M. S. S. Brooks, B. Johansson, and H. L. Skriver, in *Handbook on the Physics and Chemistry of the Actinides*, edited by A. Freeman and G. Lander (North-Holland, Amsterdam, 1984-87), vol. 1, pp. 227–269.
- [26] P. Söderlind, *Adv. Phys.* **47**, 959 (1998).
- [27] P. Söderlind, *Phys. Rev. B* **66**, 085113 (2002).
- [28] T. L. Bihan, S. Heathman, M. Idiri, G. H. Lander, J. M. Lander, J. M. Wills, A. C. Lawson, and A. Lindbaum, *Phys. Rev. B* **67**, 134102 (2003).
- [29] C.-S. Yoo, H. Cynn, and P. Söderlind, *Phys. Rev. B* **57**, 10359 (1998).
- [30] J. P. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **77**, 3865 (1996).
- [31] N. Troullier and J. L. Martins, *Phys. Rev. B* **43**, 1993 (1991).
- [32] S. Nose, *J. Chem. Phys.* **81**, 511 (1984).
- [33] G. H. Hoover, *Phys. Rev. A* **31**, 1695 (1985).
- [34] C.-S. Yoo, J. Akella, and J. A. Moriarty, *Phys. Rev. B* **48**, 15529 (1993).
- [35] J. A. Moriarty, *Phys. Rev. B* **42**, 1609 (1990).
- [36] O. J. W. ed., *Plutonium Handbook: A Guide to the Technology* (American Nuclear Society, La Grange Park, 1980).

- [37] H. L. Skriver, *The LMTO Method* (Springer, Berlin, 1984).
- [38] P. Doloffre, *Ph.D Thesis* (de l'universite Paris XI Orsay, 1997).
- [39] B. Sadigh and W. G. Wolfer, Phys. Rev. B **72**, 205122 (2005).
- [40] A. Kokalj, Comp. Mater. Sci. **28**, 155 (2003), code available from <http://www.xcrysden.org/>.