



Adaptive coordinate, real-space electronic structure calculations on parallel computers

Citation

Zumbach, Gil, N.A. Modine, and Efthimios Kaxiras. 1996. "Adaptive Coordinate, Real-Space Electronic Structure Calculations on Parallel Computers." *Solid State Communications* 99 (2): 57–61. [https://doi.org/10.1016/s0038-1098\(96\)80049-4](https://doi.org/10.1016/s0038-1098(96)80049-4).

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:41384125>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Adaptive coordinate, real-space electronic structure calculations on parallel computers

Gil Zumbach, N. A. Modine and Efthimios Kaxiras
Department of Physics, Harvard University, Cambridge MA 02138
(August 16, 2018)

We present a method for electronic structure calculations that retains all of the advantages of real space and addresses the inherent inefficiency of a regular grid, which has equal precision everywhere. The computations are carried out on a *regular* mesh in *curvilinear space*, which allows natural and efficient decomposition on parallel computers, and effective use of iterative numerical methods. A novel feature is the use of error analysis to optimize the curvilinear grid for highly inhomogeneous electronic distributions. We report accurate all-electron calculations for H_2 , O , and O_2 .

Ab initio electronic structure calculations are computationally very challenging because the singular Coulomb potential of the ions results in highly localized core wavefunctions with cusps at the ionic positions. Even when pseudopotentials are used to eliminate the core electrons, it is often desirable to treat valence electrons with highly localized wavefunctions (e.g., 1s, 2p, 3d, or 4f valence electrons), on the same footing as delocalized ones. Typical implementations use a basis in a (one-particle) Hilbert space, the choice of which requires a tradeoff between simplicity and fast convergence of physical quantities with the basis size. The simplest basis consists of plane waves. Its main drawback is uniform precision, leading to slow convergence for inhomogeneous systems like atoms, molecules, clusters, or solid surfaces. On the other hand, bases such as linearized augmented plane waves (LAPW) or muffin tin orbitals (LMTO) can be tailored to specific physical problems and therefore have excellent convergence properties. However, they lead to very complex equations. A promising alternative is a real space approach. All terms are local except for the Laplacian, which has a very short range. The resulting sparse Hamiltonian allows effective use of iterative algorithms, which vastly reduce both memory and time requirements, and is a prerequisite to any $O(N)$ treatment of electronic structure.

Harnessing the computational power of massively parallel architectures imposes additional constraints on the choice of basis. To achieve good load balance, computational complexity and memory requirements must be evenly divided among processors, a task made very difficult by complex bases like LAPW and LMTO. Another important consideration is the minimization and localization of communication between processors. Since Fourier transforms (the underlying operations in a plane wave basis) require communication between all processors, even plane waves are not an efficient basis in this respect. In contrast, a *regular* grid in real space is a very natural choice for a massively parallel computer architecture: assigning an equal section of the grid to each processor provides good load balance, minimizes interprocessor communication, and produces communication patterns that are both local and conflict-free.

Chelikowsky and collaborators [1] have reported real space electronic structure computations using a regular grid. Recently, Briggs et al. [2] have used multi-grid acceleration to improve efficiency. Yet, a regular grid in real space suffers from the same drawbacks as a plane wave basis, i.e., it has the same resolution in every region of space. Attempts to circumvent this problem have been pursued by Cho et al. [3] and by Wei et al. [4], using wavelets as a basis. Another approach investigated by Tsuchida et al. [5] uses finite-elements with a non-uniform grid. Finally, Bylaska et al. [6] have reported calculations using multi-grid methods to enhance precision locally. Irrespective of whether the formalism is based on wavelets, finite-elements, or multi-grids, enhancing the resolution by locally adding more basis elements ruins the natural mapping onto a parallel architecture.

Progress toward overcoming the limitations of plane wave bases has also been reported recently. Gygi [7] introduced the concept of adapted plane waves, a distortion of Fourier space that allows treatment of physical space with different degrees of precision. Following that development, Hamann [8] and Devenyi et al. [9] reported calculations using similar approaches. The adaptive plane wave approach eliminates the major drawback of the standard plane wave basis, but lacks the simplicity, sparseness, and natural parallelization properties of real space algorithms.

In the present work, we combine the advantages of real space calculations and those of adaptive coordinates in a scheme for performing electronic structure calculations in the context of density functional theory and the local density approximation (DFT/LDA) [10], or the generalized gradient approximation [11]. Our scheme is efficient and accurate for systems with very inhomogeneous charge distributions and takes full advantage of massively parallel computer architectures. The central idea is to work on a regular grid, but in a curvilinear space $\vec{\xi}$. The change of coordinates $\vec{x}(\vec{\xi})$, generates a *single* grid in real space \vec{x} , which is finer where high precision is needed. We refer to our method as the Adaptive Coordinate, Real-space, Electronic Structure (ACRES) algorithm. It embodies the following advantageous features:

(1) It can achieve an essentially optimal distribution of grid points. This is accomplished by a versatile choice of the curvilinear coordinates, through which a grid of fixed size is adapted to provide resolution commensurate with the physics. The only cost of the adaptation is the introduction of a nontrivial metric $g^{\alpha\beta}(\vec{\xi})$.

(2) The coordinate transform is chosen so as to minimize the discretization error. The idea is to determine an *a priori* good set of curvilinear coordinates through the use of error analysis. This differs from Gygi's original approach [7] in the adaptive plane wave scheme, where the best possible change of coordinates is found through an energy minimization.

(3) Since the communications pattern remains exactly the same as that of a regular grid in real space, highly efficient parallelization is trivially accomplished.

(4) The sparsity of the equations allows us to take advantage of iterative algorithms. This makes it possible to employ rather large grids, and consequently to investigate complex systems. The computational time scales as $N \times n_e$ with N the total number of points in the 3-dimensional grid and n_e the number of electrons in the system.

We will now describe the method in more detail. The real space coordinates $x^i(\xi^\alpha; P^m)$ depend on the curvilinear coordinates ξ^α and on some set of parameters P^m that allow us to tune the change of coordinates to a particular problem. The Jacobian of the transformation is

$$J_\alpha^i(\xi; P) = \partial x^i / \partial \xi^\alpha \quad (1)$$

with $|J| = \det J$ its determinant. The trivial metric $g^{ij} = \delta_{ij}$ in real space corresponds to the metric $g^{\alpha\beta} = J^{-1\alpha}_i J^{-1\beta}_i$ in curvilinear coordinates (summation over repeated indices implied). The Laplacian operator in curvilinear space is

$$\Delta = \frac{1}{|J|} \partial_\alpha (|J| g^{\alpha\beta} \partial_\beta), \quad (2)$$

and the integrals are transformed according to $\int d^3x = \int d^3\xi |J|$. The Coulomb potential is found by solving the Poisson equation [discretized in curvilinear coordinates by means of the Laplacian, Eq. (2)] with the sum of the electronic and nuclear charge as the source.

The equations are discretized in a box of linear size Λ_i , using a finite difference scheme on a regular grid in curvilinear space $\vec{\xi}$ with N_i points in each direction [12]. Any boundary conditions, including the phase shifts required to do multiple k -point calculations for solids, can easily be implemented in this approach. In the following, we use periodic boundary conditions.

Implementation of the method presents certain challenges due to the freedom in choosing discretization schemes. The most important ones, and the manner in which we resolved them, are discussed here briefly:

(a) Equations that are equivalent in the continuum limit are not necessarily equivalent after discretization.

For example, there exist several expressions for the Laplacian which are equivalent in the continuum limit. From physical and computational considerations, it is desirable to have a self-adjoint discretization of the Laplacian. The expression given in Eq. (2) is self-adjoint after discretization if, for a fixed pair of indices (α, β) , the finite difference operators used to represent ∂_α and ∂_β are identical.

(b) The order of the finite difference approximation for the derivatives is very important. The lowest order, two-point symmetric derivative is insufficient and does not give good results. Our experience indicates that we need to use a symmetric discretization for the derivatives with at least four points (second order).

(c) The discrete representation of the nuclear charge is equally important. For an atom with atomic number Z at position \vec{R} , the nuclear charge is $\rho(\vec{\xi}) = Z \delta(\vec{\xi}; \vec{R})$ where $\delta(\vec{\xi}; \vec{R})$ is a representation of a Dirac δ function at \vec{R} on the regular grid in $\vec{\xi}$ space. Beside the normalization condition on the δ function, an important constraint on its representation on a finite grid is that the first moment of the distribution must correspond to the location of the δ function, i.e.,

$$\int d\vec{\xi} |J| \delta(\vec{\xi}; \vec{R}) \vec{x}(\vec{\xi}) = \vec{R} \quad (3)$$

We found the most useful representation to be a Gaussian

$$\delta(\vec{\xi}; \vec{R}) \propto \exp\left(-|\vec{\xi} - \vec{\xi}_0|^2 / 2\sigma^2 \Delta\xi^2\right) \quad (4)$$

with $\Delta\xi$ the regular grid spacing, σ an adjustable parameter, and $\vec{\xi}_0$ chosen to satisfy the constraint on the first moment of the distribution. This choice reduces the translational invariance problems discussed in the next point.

(d) The presence of the grid breaks translational invariance. We call the distance between an atomic center and the nearest grid point the *offset*. The energy depends on the offset, and this effect can be quite large due to the Coulomb singularity. The dependence is reduced by strong adaptation, which makes the cell of the real space grid very small near the atomic sites. The Gaussian representation of the δ function further reduces the dependence of the energy on the offset because it results in a smoother transfer of charge as the position of an atom changes. The combined use of strong adaptation and a Gaussian δ function eliminates the translation invariance problem.

(e) A final challenge is the actual choice of curvilinear coordinates. A necessary condition for the mapping between \vec{x} and $\vec{\xi}$ is that it must be one to one, i.e., the grid in x space must not be folded. As the Laplacian involves the derivative of the metric, and the metric is computed from the Jacobian, the mapping must be at least C^2 on the torus in order to ensure smoothness. It is also desirable that the mapping be spherically symmetric around an atom. We use a two level coordinate transformation $\vec{x}(\vec{\xi}; P)$ with a *global backdrop* useful for

simulating isolated structures, and further *local adaptation* around each atom position. The global backprop is a simple independent transformation along each axis $x^i = x^i(\xi^i)$ and creates a flat central region with a high density of grid points and a surrounding region with a decreasing density of grid points. The local adaptation creates a spherical deformation of the grid around each atomic center \vec{R}_ν , with the amount of adaptation A_ν and the size of the adapted region ρ_ν as independent variables. As suggested by Gygi [7], for a given $\det J(\vec{R}_\nu)$ and ρ_ν , the final results are not very sensitive to the details of the formula for $\vec{x}(\vec{\xi})$. The computations presented below were carried out with the simple form for the local adaptation

$$\vec{x}(\vec{\xi}; P) = \vec{\xi} - \sum_\nu A_\nu (\vec{\xi} - \vec{R}_\nu) \exp\left(-\frac{|\vec{\xi} - \vec{R}_\nu|^2}{2\sigma^2(A_\nu, \rho_\nu)}\right) \quad (5)$$

with the function $\sigma(A, \rho)$ chosen such that ρ gives the real space width of the adapted region.

A question of central importance is how to choose the different parameters of the grid so as to generate a nearly optimal mesh for a given physical problem. We resolve this issue by constructing an estimate for the error in the integrals and then choosing the parameters that minimize the error. To illustrate this point, consider a periodic, one-dimensional integral $I(f) = \int d\xi f(\xi)$ computed numerically on a regular mesh

$$I_N(f) = \sum_i \Delta\xi f(\xi_i) \quad (6)$$

with $\Delta\xi = \Lambda/N$. We evaluate the elementary error by comparing the integrals computed with N and $N/2$ points [13]. More precisely, with $N/2$ points, the rectangular element of integration is $\delta I_{N/2} = 2\Delta\xi f(\xi_i)$. In comparison, the same element of integration computed with N points is $\delta I_N = \Delta\xi [f(\xi_{i-1})/2 + f(\xi_i) + f(\xi_{i+1})/2]$. An elementary estimate of the error is given by

$$\delta e(f) = \delta I_N - \delta I_{N/2} = \Delta\xi^3 f_i''/2. \quad (7)$$

If the constant 1/2 on the right hand side is replaced by 1/12 we obtain a rigorous upper bound due to Peano [14]. The error in the numerical integral is then estimated by the L_2 norm of δe

$$e(f) = \frac{1}{2} \left(\frac{\Lambda}{N}\right)^{5/2} \left(\sum_i \Delta\xi (f_i'')^2\right)^{1/2}. \quad (8)$$

The above idea is easily generalizable to three-dimensional integrals (whereas the rigorous Peano bound is difficult to extend to higher dimensions). The last step is to pick an integrand f so as to obtain an *a priori* estimate of the optimal grid parameters by minimizing $e(f)$. By experimenting on several atoms we have found that

$f = |J|\rho V_{K-S}$ provides an adequate indicator. Due to large cancellations forced by the Kohn-Sham eigenvalue equation, this term gives the leading factor for the error in the total energy.

Using the approach described in this paper, we have implemented DFT/LDA [15] and DFT/GGA [11] electronic structure calculations on the CM-5 massively parallel supercomputer. Within this approach, all-electron computations involving atoms in the first row of the periodic table are feasible. We have also implemented the pseudopotential approach, using the norm-conserving nonlocal pseudopotentials of Bachelet et al. [16], and the Kleinman-Bylander procedure to render the nonlocal components separable [17].

For the all-electron calculations, the adaptation of the grid is determined by the requirement that the density of grid points near the atomic cores is sufficient to accurately represent the $1/r$ divergence of the Coulomb potential. For example, Fig. 1 shows a grid used for the H_2 molecule calculation. This clearly indicates the very large difference between the spacing of grid points in the unoccupied vacuum region and near the atomic nuclei. Fig. 2 shows the occupied wave functions of the O_2 molecule along a line through the centers of the two atoms. The enhancement of the grid resolution throughout the regions where the electronic wave functions are large and the very strong enhancement close to the nuclei allow accurate representation of the smooth tails of the wave functions as well as the cusps and nodes near the nuclei.

For a more quantitative comparison to other theoretical results and to experiment, Table I shows our calculated results for H_2 and O. It is clear from this comparison that our results are in complete agreement with other theoretical work using similar methods. Therefore, the difference between calculated values and experimental measurements reflects fundamental limitations of the underlying theory (DFT/LDA or DFT/GGA), rather than limitations in the accuracy of our method.

The authors are grateful to F. Gygi for sharing his insight on adaptive methods. This work was supported by the Office of Naval Research grant N00014-93-1-0190.

-
- [1] J. R. Chelikowsky, N. Troullier, and Y. Saad, Phys. Rev. Lett. **72**, 1240 (1994); J. R. Chelikowsky, N. Troullier, K. Wu, and Y. Saad, Phys. Rev. B **50**, 11355 (1994).
 - [2] E. L. Briggs, D. J. Sullivan, and J. Bernholc, preprint.
 - [3] K. Cho, T. A. Arias, J. D. Joannopoulos, and P. K. Lam, Phys. Rev. Lett. **71**, 1808 (1993).
 - [4] S. Q. Wei and M. Y. Chou, Bull. Amer. Phys. Soc. **40**, 43 (1995).
 - [5] E. Tsuchida and M. Tsukada, Sol. St. Comm. **94**, 5 (1995).
 - [6] E. J. Bylaska *et al.*, preprint.

- [7] F. Gygi, Europhys. Lett. **19**, 617 (1992); Phys. Rev. B **48**, 11692 (1993); **51**, 11190 (1995); Private communication.
- [8] D. R. Hamann, Phys. Rev. B **51**, 7337 (1995); **51**, 9508 (1995).
- [9] A. Devenyi, K. Cho, T. A. Arias, and J. D. Joannopoulos, Phys. Rev. B **49**, 13373 (1994).
- [10] P. Hohenberg and W. Kohn, Phys. Rev. **136**, B864 (1964); W. Kohn and L. Sham, *ibid.* **140**, A1133 (1965).
- [11] J. P. Perdew, in *Electronic Structure of Solids '91*, edited by P. Ziesche and H. Eschrig (Akademie Verlag, Berlin, 1991).
- [12] Through our finite difference scheme, we have an approximation of the original equations, rather than a projection of the original problem onto a basis. Consequently, the variational aspect of a basis is lost, and our electronic eigenvalues are not necessarily upper bounds of the true electronic eigenvalues.
- [13] We caution the reader against a subtle but important pitfall in the error estimation of a numerical integral: the second Euler-Maclaurin summation formula is only an asymptotic formula. When applied to the case of a discretized integral of a C^∞ periodic function, it leads to an obvious contradiction!
- [14] P. J. Davis and P. Rabinowitz, *Methods of Numerical Integration*, 2nd ed. (Academic Press, Orlando, FL, 1984).
- [15] J. P. Perdew and A. Zunger, Phys. Rev. B **23**, 5048 (1981).
- [16] G. B. Bachelet, D. R. Hamann, and M. Schlüter, Phys. Rev. B **26**, 4199 (1982).
- [17] L. Kleinman and D. M. Bylander, Phys. Rev. Lett. **48**, 1425 (1982).
- [18] B. G. Johnson, P. M. W. Gill, and J. A. Pople, J. Chem. Phys. **98**, 5612 (1992).
- [19] K. P. Huber and G. Herzberg, *Molecular Spectra and Molecular Structure* (Van Nostrand Reinhold Company, New York, 1979), Vol. IV.
- [20] J. P. Perdew *et al.*, Phys. Rev. B **46**, 6671 (1992).
- [21] Y.-M. Juan and E. Kaxiras, Phys. Rev. B **48**, 14944 (1993); Private communication.
- [22] C. E. Moore, *Atomic Energy Levels* (U. S. Government Printing Office, Washington, 1971), Vol. I.

FIG. 1. The $24 \times 12 \times 12$ [a.u.] grid used for H_2 , in a horizontal cross-section through the atoms (every fourth line shown). Notice the effect of the global backdrop (crosslike pattern) and the local adaptation around each atom.

FIG. 2. Occupied wave functions of the O_2 molecule, along a line through the centers of the atoms. The π bonding and anti-bonding wave functions collapse onto the horizontal axis (they have nodes through the atomic centers). The $1s$ bonding and anti-bonding states were scaled by a factor of $1/3$ so they could be displayed on the same scale. Points on the curves indicate values at actual grid points used in the calculation.

TABLE I. Calculated bond length a_0 , vibrational frequency ω , and minimum energy E_0 for H_2 and atomic energy E_{at} of O. The zero-point vibrational energy has been subtracted from the experimental total energy of H_2 .

	ACRES	Other DFT Theory	Experiment
H_2 [LDA]			
a_0 (a.u.)	1.448	1.446 ^a	1.401 ^b
ω (cm ⁻¹)	4192	4207 ^a	4401 ^b
E_0 (Ry)	-2.276	-2.27 ^c	-2.349 ^b
H_2 [GGA]			
a_0 (a.u.)	1.416	1.413 ^a	1.401 ^b
ω (cm ⁻¹)	4381	4373 ^a	4401 ^b
E_0 (Ry)	-2.340	-2.34 ^c	-2.349 ^b
O [LDA]			
E_{at} (Ry)	-148.870	-148.938 ^d	-150.027 ^e
O [GGA]			
E_{at} (Ry)	-149.912	-149.994 ^d	-150.027 ^e

^a Reference [18], S-VWN and B-LYP

^b Reference [19]

^c Reference [20], LSD and PW GGA-II

^d Reference [21], LDA and PW91

^e Reference [22], Spin unpolarized ground state, $2p^4 \ ^1D$



