

Multigrid method for electronic structure calculations

M. Heiskanen, T. Torsti, M. J. Puska, and R. M. Nieminen

Laboratory of Physics, Helsinki University of Technology, P.O. Box 1100, FIN-02015 HUT, Finland

(Received 10 January 2001; published 1 June 2001)

A general real-space multigrid algorithm for the self-consistent solution of the Kohn-Sham equations appearing in the state-of-the-art electronic-structure calculations is described. The most important part of the method is the multigrid solver for the Schrödinger equation. Our choice is the Rayleigh quotient multigrid method (RQMG), which applies directly to the minimization of the Rayleigh quotient on the finest level. Very coarse correction grids can be used, because there is, in principle, no need to represent the states on the coarse levels. The RQMG method is generalized for the simultaneous solution of all the states of the system using a penalty functional to keep the states orthogonal. The performance of the scheme is demonstrated by applying it in a few molecular and solid-state systems described by nonlocal norm-conserving pseudopotentials.

DOI: 10.1103/PhysRevB.63.245106

PACS number(s): 71.15.Dx, 31.15.Ew

I. INTRODUCTION

One of the goals of computational materials science is to calculate from first principles the various physical and chemical properties. This requires the solution of the electronic and ionic structures of the materials system in question. The density-functional theory (DFT) makes a huge step towards this goal by casting the untractable problem of many interacting electrons to that of noninteracting particles under the influence of an effective potential.¹ The adiabatic approximation allows one to separate the ionic degrees of freedom from those of the electrons. However, in order to apply DFT in practice one has to resort to approximations for electron exchange and correlation such as the local-density approximation (LDA) or the generalized-gradient approximation. Moreover, in the case of systems consisting of hundreds or more atoms it is still a challenge to solve numerically efficiently for the ensuing Kohn-Sham equations.

The numerical solution of the Kohn-Sham equations is the concern of our present paper. It deals with real-space (RS) methods, in which the values of the different functions are presented using three-dimensional point grids, and the partial differential equations are discretized using finite differences.^{2,3} The RS methods, as suggested by the name chosen, are contrasted with the popular plane-wave (PW) schemes.^{4,5} There are several aspects favoring the RS methods over the PW methods. Both of the methods are used in the context of pseudopotentials describing the electron-ion interactions, but only the RS can easily be used in all-electron calculations or with hard pseudopotentials of, i.e., first-row or transition-metal atoms, because the RS grid can be refined in a natural way in the ion core regions (composite grids,⁶⁻⁸ adaptive coordinates⁹⁻¹¹). Systems, such as surfaces, containing different length scales are more economically described in the RS than in the PW scheme because one need not waste many grid points in the vacuum regions to describe the slowly varying tails of wave functions. In the RS methods, periodic boundary conditions are not necessary. This leads to ease and accuracy in describing charged atomic clusters in contrast to PW methods, which require an artificial neutralizing background charge. Besides the above ‘‘physical arguments,’’ there are also methodological and computational aspects favoring the RS methods. The RS methods allow a systematic convergence control by increas-

ing the grid (or basis function) density. (The PW methods do so also, by adjusting the cutoff energy of the plane-wave expansions.) The so-called ‘‘order- N ’’ methods,¹² the computational cost of what scales linearly with the number of electrons, require localized real-space wave functions leading naturally to the employment of RS methods.¹³ The discretizations in the real-space grid can be made local, and therefore, parallelization can effectively use data decomposition in which different real-space regions are handled with different processing units and the communications between processing units will be mainly short ranged.¹⁴

More specifically, our choice for the numerical method is a multigrid scheme.^{15,2} Several approaches employing the multigrid idea within electronic structure calculations have appeared during recent years.^{16,14,17,8,18} The main idea of multigrid methods is that they avoid the critical slowing-down (CSD) phenomenon occurring when a partial differential equation discretized on a real-space grid is solved with a simple relaxation method such as the Gauss-Seidel method. The discretization operators typically use information from a rather localized region of the grid at a time. Therefore, the high-frequency error of the length scale of the grid spacing is reduced very rapidly in the relaxation. However, once the high-frequency error has effectively been removed, the very slow convergence of the low-frequency components dominates the overall error reduction rate,¹⁵ i.e., CSD occurs. In multigrid methods, one stops the relaxation on a given (fine) grid before CSD sets in and transfers the equation to a coarser grid (the so-called restriction operation) where the low-frequency components can be solved more efficiently. On the coarsest grid, the problem is solved exactly or as accurately as possible, after which, one interpolates (the so-called prolongation operation) the correction to finer grids, performing simultaneously relaxations in order to remove the high-frequency errors introduced in the interpolation.

The solution of the Poisson equation by multigrid methods is straightforward.¹⁵ This is because the error (or the correction needed) also obeys a Poisson equation, and thus, will be a smooth function to be presented and solved on the repeatedly coarser grids optimal to handle the lower frequencies. The solution of an eigenvalue problem, such as the Schrödinger equation, is a much more complicated task than that of the Poisson equation. The problem is no more linear

because both the eigenfunction and the eigenvalue have to be solved simultaneously; then the error no longer obeys the same equation as the solution. Also, one has to solve for several eigenpairs (eigenvalues and corresponding eigenvectors). Moreover, the existence of both negative and positive eigenvalues makes the problem indefinite. This implies severe difficulties for many simple iterative methods that converge only in the case of a positive definite iteration matrix. In particular, it can easily be shown that when using Gauss-Seidel relaxation for the Schrödinger equation, the high-frequency components typically converge, as in the case of the Poisson equation, but the low-frequency components may diverge, although the divergence may be slow.¹⁹ More complicated methods, such as Kaczmarz relaxation, are guaranteed to converge, but may have clearly inferior high-frequency reduction rates, which are essential for the overall speed of multigrid methods. Another convergent method is the generalized minimal residual algorithm (GMRES), which is considerably more complex than Gauss-Seidel relaxation.²⁰

A standard recipe for dealing with eigenproblems with multigrids is the full-approximation-storage (FAS) method originally described by Brandt *et al.*²¹ In FAS, one solves for the entire problem on the coarse grids also and ends up solving for a properly modified problem so that its solution can be used in correcting the fine-grid solution. The FAS method may not be very straightforward to implement the Schrödinger equation. It is also difficult to present some actual potential on the coarse levels accurately enough. However, some successful applications of FAS have appeared in the context of electronic structure calculations by Beck *et al.*²² and Wang and Beck¹⁸ and advanced strategies for FAS have been proposed.²³

Briggs *et al.*^{16,14} employ a multigrid method in electronic structure calculations by linearizing the Schrödinger problem and presenting the potential contribution on the coarse levels by an error term (residual) only. Then, on the coarse levels, they solve effectively for the Poisson problem. Ancilotto *et al.*¹⁷ modified the method by Briggs *et al.* by shifting to a full multigrid (FMG) scheme and by solving on the coarse grids, a problem including a local potential term. The idea of FMG is to start the smoothing iterations from a coarse grid. Then the interpolation to a finer grid provides a good initial guess of the solution. The FMG scheme can accelerate the convergence remarkably with respect to the (above-described) V-cycle scheme in which one starts from the finest level. Fattebert⁸ used a multigrid method with a block Galerkin inverse iteration (BGII) and GMRES in the relaxations. In the method, the current approximation is kept orthogonal against all the nearby states during the multigrid cycle. The inverse iteration converges for a given guess for the energy eigenvalue towards the nearest eigenvalue. In order to solve all the desired lowest eigenstates a good guess for the eigenvalue spectrum is needed in the beginning of iterations, but thereafter, large computational savings may be expected because explicit orthogonalizations are not needed (at least not between well-separated states).

A severe problem in the existing multigrid schemes for the Schrödinger equation is often that the coarse grids cannot

well approximate the solutions of the coarse grid equations themselves. As a consequence, the correction from coarse grids, no matter how accurately the equation is solved, may be ineffective in correcting the fine grid solution, and as a result, the overall process converges slowly. Therefore, one is restricted to the use of rather fine grids only and the convergence speed of the scheme is drastically lowered. In those multigrid methods, that use the potential on the coarse grids, the size of the coarsest grid has been typically of the order of $31 \times 31 \times 31$.^{8,17} However, using the FAS method, coarser grids are possible at least for systems with a small number of eigenstates solved even when the all-electron scheme is employed.¹⁸ If a large number of eigenstates have to be solved, problems may arise, because the coarse grids may not be able to represent eigenstates with many nodes or the ordering of the states may change between the successive grids. To bypass these problems in FAS, rather complicated strategies are needed.²³

In order to avoid the coarse grid representation problems, we utilize the so-called Rayleigh Quotient Multigrid (RQMG) method introduced by Mandel and McCormick.²⁴ In this method, the coarse grid relaxation passes are performed so that the Rayleigh quotient calculated on the *fine* grid will be minimized. In this way, there is no requirement for the solution to be well represented on a coarse-grid and the coarse-grid representation problem is avoided. Mandel and McCormick²⁴ introduced the method for the solution of the eigenpair corresponding to the lowest eigenvalue. We have generalized it to the simultaneous solution of a desired number of lowest eigenenergy states by developing a scheme that keeps the eigenstates separated by the use of a penalty functional, Gram-Schmidt orthogonalization, and subspace rotations. Our generalization of the RQMG method is an attractive alternative for large-scale electronic structure calculations.

The Kohn-Sham equations have to be solved self consistently, i.e., the wave functions solved from the single-particle equation determine via the density (solution of the Poisson equation and the calculation of the exchange-correlation potential) the effective potential for which they should again be solved. To approach this self consistency requires an optimized strategy so that numerical accuracy of the wave functions and the potential increase in balance, enabling the most efficient convergence.¹⁸ In order to avoid the divergence of the self-consistency iterations, the mixing of the input and output solutions is needed. For this feedback procedure, sophisticated schemes²⁵ and control strategies¹¹ have been presented.

The outline of the present paper is as follows. In Sec. II we represent the most important ideas of the density-functional theory. Section III is devoted for numerical methods, the most important of which is the Schrödinger equation solver developed; the strategy for the self-consistency iterations is also discussed. In Sec. IV we demonstrate with the help of a couple of examples, the performance of our scheme in calculating the electronic structures of small molecules and solid-state systems described by pseudopotentials. Section V summarizes the work and gives outlines for the future developments.

II. THE KOHN-SHAM SCHEME

In the Kohn-Sham method for electronic structure calculations, one solves for a set of equations self consistently.¹ In the following, we present the equations in the spin-compensated form. In practice, we have made the straightforward generalization using the spin-density functional theory. The set of equations reads as (atomic units with $\hbar = m_e = e = 1$ are used):

$$\left(-\frac{1}{2}\nabla^2 + V_{\text{eff}}(\mathbf{r})\right)\Psi_i = \epsilon_i\Psi_i, \quad (1)$$

$$n(\mathbf{r}) = \sum_i^N |\Psi_i(\mathbf{r})|^2, \quad (2)$$

$$V_{\text{eff}}(\mathbf{r}) = V_{\text{ion}}(\mathbf{r}) + V_{\text{H}}(\mathbf{r}) + V_{\text{XC}}(\mathbf{r}), \quad (3)$$

$$V_{\text{H}}(\mathbf{r}) = \int \frac{n(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}', \quad (4)$$

$$V_{\text{XC}}(\mathbf{r}) = \frac{\delta E_{\text{XC}}[n(\mathbf{r})]}{\delta n(\mathbf{r})}. \quad (5)$$

The first equation (1) is a Schrödinger equation for noninteracting particles in an effective potential $V_{\text{eff}}(\mathbf{r})$. For finite systems, the wave functions are required to vanish at the boundaries of the computation volume. In the case of infinite periodic systems, the complex wave functions have to obey the Bloch theorem at the cell boundaries. The electron density $n(\mathbf{r})$ is obtained from a sum over the N occupied states. The effective potential consists of an external potential $V_{\text{ion}}(\mathbf{r})$ due to ions (or nuclei in all-electron calculations), the Hartree potential $V_{\text{H}}(\mathbf{r})$ calculated from the electron-density distribution, and the exchange-correlation potential $V_{\text{XC}}(\mathbf{r})$. In the examples of the present paper, we use the norm-conserving non-local pseudopotentials for the electron-ion interactions and the local-density approximation (LDA) for the exchange-correlation energy

$$E_{\text{XC}}[n] = \int \epsilon_{\text{XC}}(n(\mathbf{r}))n(\mathbf{r})d\mathbf{r}, \quad (6)$$

and for the exchange-correlation potential

$$V_{\text{XC}}(\mathbf{r}) = \epsilon_{\text{XC}}(n(\mathbf{r})) + n(\mathbf{r}) \left. \frac{d\epsilon_{\text{XC}}}{dn} \right|_{n=n(\mathbf{r})}. \quad (7)$$

The Hartree potential is solved from the Poisson equation

$$\nabla^2 V_{\text{H}}(\mathbf{r}) = -4\pi n(\mathbf{r}). \quad (8)$$

In practice, the electron-density $n(\mathbf{r})$ is substituted by the total charge-density $\rho(\mathbf{r})$, which includes the positive ionic (nuclear) charge neutralizing the system. In the case of finite systems, Dirichlet boundary conditions are used with the Coulomb potential values calculated using a multipole expansion. For periodic systems, we fix the average Coulomb potential to zero and allow the periodic boundary conditions to result in the corresponding converged potential.

The self-consistent solution of the above Kohn-Sham equations leads to the ground state electronic structure minimizing the total energy

$$E_{\text{tot}} = \sum_i \int \Psi_i^*(\mathbf{r}) \left(-\frac{1}{2}\nabla^2 \right) \Psi_i(\mathbf{r}) d\mathbf{r} + \frac{1}{2} \int V_{\text{H}}(\mathbf{r})n(\mathbf{r})d\mathbf{r} + \int V_{\text{ion}}(\mathbf{r})n(\mathbf{r})d\mathbf{r} + E_{\text{XC}} + E_{\text{ion-ion}}, \quad (9)$$

where $E_{\text{ion-ion}}$ is the repulsive interaction between the ions (nuclei) of the system. Instead of the self-consistency iterations, the solution of the Kohn-Sham problem can be found by minimizing directly the total energy with respect to the wave-function parameters, e.g., plane-wave coefficients.⁴ However, Kresse and Furthmüller^{5,25} have found this scheme less efficient than the self-consistency iterations.

III. NUMERICAL METHODS

A. Schrödinger equation solver

In our real-space method, we start from an initial guess for the effective potential and initial wave functions generated by random numbers in grid points. The wave functions and the Hartree potential are updated alternately towards self consistency. The solution of the Poisson equation is a standard task for the multigrid scheme. If a reasonable guess for the Coulomb potential is not available, the FMG method will produce the solution starting from random numbers and requiring the work that scales linearly as a function of the size of the system [$O(N)$]. During the Kohn-Sham iterations, one can start from the present approximation of Coulomb potential and update it with respect to the new charge density by performing only a few V cycles.

The solution of the wave functions is a much more complicated task than that of the Poisson equation because one has to solve an eigenvalue problem that in the state-of-the-art electronic structure calculations means the determination of several hundreds of eigenpairs. For this purpose, we have developed a scheme based on RQMG method introduced by Mandel and McCormick²⁴ for the solution of the eigenpair corresponding to the lowest eigenvalue. We begin by reviewing the basic principles of RQMG. This is most easily done in the framework of the so-called coordinate relaxation method. Thereafter, we go through the modifications made in order to simultaneously solve for several eigenpairs.

Coordinate relaxation is a method of solving the discretized eigenproblem

$$Hu = \lambda Bu \quad (10)$$

by minimizing the Rayleigh quotient

$$\frac{\langle u|H|u \rangle}{\langle u|B|u \rangle}. \quad (11)$$

Above, H and B are matrix operators chosen so that the Schrödinger equation discretized on a real-space point grid with spacing h is satisfied to a chosen order $O(h^n)$. In Eq. (11) u is a vector containing the wave-function values at the

grid points. In the relaxation method, the current estimate u is replaced by itself plus a multiple of some search-vector d

$$u' = u + \alpha d, \quad (12)$$

and α is chosen to minimize the Rayleigh quotient. This leads to a simple quadratic equation for α . [Find the minimum of the expression (14) below with respect to α . In the case of a complex wave function, one has to solve for the real and imaginary parts of α from a coupled pair of quadratic equations.] Moreover, the search-vector d is simply chosen to be unity in one grid point and to vanish in all other points. A complete coordinate relaxation pass is then obtained by performing the minimization at each point in turn and these passes can be repeated until the lowest state is found with desired accuracy.

Naturally, the coordinate relaxation also suffers from CSD because of the use of local information only in updating u in a certain point. In order to avoid it, one applies the multigrid idea. In the multigrid scheme by Mandel and McCormick,²⁴ the crucial point is that *coarse-grid* coordinate relaxation passes are performed so that the Rayleigh quotient calculated on the *fine* grid will be minimized. In this way, there is no requirement for the solution to be well represented on a coarse grid. In practice, a coarse-grid search substitutes the fine-grid solution by

$$u'_f = u_f + \alpha I_c^f d_c, \quad (13)$$

where the subscripts f and c stand for the fine and coarse grids, respectively, and I_c^f a prolongation operator interpolating the coarse grid vector to the fine grid. The Rayleigh quotient to be minimized is then

$$\begin{aligned} & \frac{\langle u_f + \alpha I_c^f d_c | H_f | u_f + \alpha I_c^f d_c \rangle}{\langle u_f + \alpha I_c^f d_c | B_f | u_f + \alpha I_c^f d_c \rangle} \\ &= \frac{\langle u_f | H_f u_f \rangle + 2\alpha \langle I_c^f H_f u_f | d_c \rangle + \alpha^2 \langle d_c | H_c d_c \rangle}{\langle u_f | B_f u_f \rangle + 2\alpha \langle I_c^f B_f u_f | d_c \rangle + \alpha^2 \langle d_c | B_c d_c \rangle}. \end{aligned} \quad (14)$$

The second form is obtained by relating the coarse-grid operators H_c and B_c with the fine-grid ones, H_f and B_f by the Galerkin condition

$$H_c = I_f^c H_f I_c^f, \quad (15)$$

$$B_c = I_f^c B_f I_c^f,$$

and the restriction operator I_f^c has to be the transpose of the prolongation operator

$$I_f^c = (I_c^f)^T. \quad (16)$$

The key point to note is that when $H_f u_f$ and $B_f u_f$ are provided from the fine grid to the coarse grid, the remaining integrals can be calculated on the coarse grid itself. Thus, one really applies coordinate relaxation on the coarse grids to minimize the *fine level* Rayleigh quotient. This is a major departure from the earlier methods, which to some extent, rely on the ability to represent the solution of some coarse-

grid equation on the coarse grid itself. Here, on the other hand, one can calculate the *exact* change in the Rayleigh quotient due to *any* coarse-grid change, no matter how coarse the grid itself is. There is no equation whose solution would have to be representable.

Thus, in the Rayleigh quotient minimization multigrid (RQMG) algorithm, the coordinate relaxation passes on each level keep track of the integrals in Eq. (14). Actually, on the finest level, we use Gauss-Seidel relaxation, which very effectively smooths the errors of the wave length corresponding to the grid spacing. When calculating several eigenpairs, Gauss-Seidel relaxation may also work as a *residual minimization* method. The idea is that the coarse-grid iterations with Gram-Schmidt orthogonalization can provide the separation of the eigenstates so well that the subsequent finest level relaxations converge to the correct (nearest) eigenstates without orthogonalization. This requires that the effect of the coarse-level smoothings on the low-frequency components of the solutions overcomes the possible divergence tendency of these components caused by the Gauss-Seidel relaxation on the finest level.

Moreover, we discretize the original equation separately on each grid [discretization coarse grid approximation (DCA)] instead of using the Galerkin conditions of Eq. (15). This may in principle decrease the convergence rate and force a limit to the coarsest possible grid in order to avoid instability or divergence. However, we have observed this DCA implementation of RQMG to be quite stable and efficient. To avoid possible coarse-level instabilities occurring, especially during the first few iteration cycles, we may recalculate the Rayleigh quotient whenever coarse grid corrections are interpolated to a finer grid. Later, when approaching the convergence, the recalculation can be omitted.

For the matrix operators H and B we have used either high-order [$O(h^4)$ or higher] Mehrstellen or central difference stencils (CDS).^{14,8} The use of high-order stencils reduces remarkably the density of grid points needed. The benefit of the Mehrstellen scheme is that more local information is used. The scheme leads to controlled accuracy and convergence properties and to more isotropic smoothing of the error in comparison with the use of CDS's. The local nature also enables a more efficient parallel coding. As the prolongation operator I_c^f we usually use trilinear interpolation and as the restriction operator I_f^c its transpose, the so-called full-weighting operator, in which the coarse-grid values are chosen to be the averaged values of the surrounding fine-grid points. The integrations are performed by the trapezoidal rule.

Next we consider the generalization of the RQMG method to the simultaneous solution of several (N) mutually orthogonal eigenpairs. The separation of the different states is divided into two or three subtasks. First, in order to make the coarse-grid relaxations converge towards the desired state, we apply a penalty functional scheme. Given the k lowest eigenfunctions, the next lowest, $(k+1)$ th state is searched for by minimizing the functional

$$\frac{\langle u_{k+1} | H | u_{k+1} \rangle}{\langle u_{k+1} | B | u_{k+1} \rangle} + \sum_{i=1}^k q_i \frac{\langle u_i | u_{k+1} \rangle^2}{\langle u_i | u_i \rangle \cdot \langle u_{k+1} | u_{k+1} \rangle}. \quad (17)$$

The overlap integral in the penalty term is squared to make the penalty positive definite. The denominator is required to make the functional independent of the norms of u_i , $i = 1, \dots, k+1$. The minimization of this functional is equivalent to imposing the orthonormality constraints against the lower k states, when $q_i \rightarrow \infty$. By increasing the shifts q_i , any desired accuracy can be obtained, but in order to obtain a computationally efficient algorithm, a reasonable finite value should be used, for example

$$q_i = (\lambda_{k+1} - \lambda_i) + Q, \quad (18)$$

where Q is a sufficiently large positive constant. In our test calculations Q is of the order of $Q = 0.5, \dots, 2$ Ha.

We minimize the expression (17) *simultaneously* for all N states. This simplifies the algorithm and enables a future parallelization over the eigenstates. Thus, the current approximations are used for u_i , $i = 1, \dots, k$. Moreover, changes in the u_i during a given relaxation sweep are not used to update the penalty term in Eq. (17). This is sufficient, when the states are always ordered in the same way, in the order of increasing eigenvalue. In order to reduce computations, the B innerproduct is actually used in calculating the penalty term integrals because the values of Bu are readily available from the finer level. The substitution (13) is introduced in the functional (17) and the minimization with respect to α leads again to a quadratic equation. This time the coefficients contain terms due to the penalty part.

On the finest level, we do not apply the minimization of the penalty functional. The ideal situation would be if a residual minimization method, such as the Gauss-Seidel method, would keep the states calculated on the coarse levels separated. We found out in practical calculations that this is not true, at least when the states are far from convergence. Therefore, we have developed for the finest level a scheme, which by employing Gram-Schmidt orthogonalization and subspace rotation, keeps the eigenstates orthogonal. The subspace rotation is a method to find the most optimally separated eigenvectors from the approximative ones. The major steps of the rotation are:

(i) Calculation of the Hamiltonian matrix elements between the current states:

$$\bar{H}_{i,j} = \langle u_i | B^{-1} H | u_j \rangle. \quad (19)$$

(ii) Calculation of the overlap matrix:

$$\bar{S}_{i,j} = \langle u_i | u_j \rangle. \quad (20)$$

The use of matrix elements of Eqs. (19) and (20) leads to eigenvectors orthogonal in the desired Euclidian sense (I orthogonal) and not in the sense of the B innerproduct.

(iii) Diagonalization to find the optimal eigenvectors ($u'_k = \sum_j \bar{A}_{k,j} u_j$) and corresponding eigenvalues (λ_k):

$$\sum_j \bar{H}_{i,j} \bar{A}_{j,k} = \lambda_k \sum_j \bar{S}_{i,j} \bar{A}_{j,k}. \quad (21)$$

In practice, we apply the approximation

$$\langle u_i | B^{-1} H | u_j \rangle \approx \langle u_i | u_j \rangle \frac{\langle u_i | H u_j \rangle}{\langle u_i | B u_j \rangle}. \quad (22)$$

The Gram-Schmidt orthogonalization and the subspace rotations are organized so that the space of the eigenvectors is first divided into small clusters corresponding to close eigenvalues. The Gram-Schmidt orthogonalization is then performed for each cluster at a time so that its eigenvectors become orthogonal against the eigenvectors of the clusters of lower eigenvalues. Then a subspace rotation is performed within the states belonging to the present cluster. The division to clusters reduces remarkably the cost of the subspace rotation. This is because the cost is proportional to $O(N^3)$, where N is the number of states rotated. Moreover, the subspace rotation requires the calculation of matrix elements that are more complicated than those for the simple Gram-Schmidt orthogonalization.

According to our test calculations, this subspace rotation scheme leads quite effectively to I -orthogonal eigenstates. This is seen as a convergence of the eigenvalue problem within the numerical accuracy, i.e., the residuals of different eigenstates vanish. However, in order to achieve exactly vanishing residuals, very accurate eigenvalues are needed—the residual norm scales as the square root of the error in the eigenvalue. Therefore one introduces for every state a constant additional potential equal to the current eigenvalue estimate, which shifts the eigenvalue towards zero. This effectively increases the number of significant digits that can be obtained. The error in the eigenvalue scales as the square of the residual. When applying the subspace rotation it is important to complete the highest eigenvalue cluster; otherwise the rotation may become inefficient.

The orthogonalization needed scales as $O(N^3)$. For small systems of several tens of eigenpairs this is not yet a problem. The algorithm is effective and the number of fine-grid orthogonalizations remains quite plausible, for example, in comparison with the conjugate gradient search of eigenpairs employing only the finest grid.²⁶ But for larger systems with hundreds of states it will be the bottleneck. One solution could be to rely on the finest level only on a residual minimization method when the initial stages of the iteration process have been performed and the solution is clearly on a stable track towards convergence.

B. Strategy for self consistency iterations

The Kohn-Sham problem has to be solved self-consistently. This means that an optimal strategy is needed so that computing time is not wasted in the beginning of the self-consistency iterations to obtain unnecessarily accurate wave functions, because these will change during the later iterations due to the changes in the potential. Updating the potential, including the solution of the Poisson equation, is a much less time-consuming task than the update of all the wave functions. Therefore, the potential update can be performed frequently.¹⁸

The examples of this paper are small-molecule and bulk-solid systems described by pseudopotentials. The strategy used is schematically presented in Fig. 1. Similar strategies




FIG. 1. Strategy of self-consistency iterations. First, the wave functions are solved nonself-consistently using the full multigrid method in the initial potential corresponding to the superposition of pseudoatoms. Then the effective potential is updated (this is denoted by P in the figure). The potential update amounts to calculation of the new electron density, the solution of the Poisson equation, and calculation of the new exchange correlation potential. Next, the wave functions are updated by one V cycle. These two steps are repeated until self-consistency has been reached.

can certainly be applied in other kinds of Kohn-Sham calculations, for example, in those employing all-electron or jellium-type models. In the examples of this paper the initial electron density is the superposition of the pseudoatom densities centered around given nuclear positions. From the superposition, we calculate the initial effective potential, where the wave functions are solved accurately enough using the full-multigrid method. The FMG process is started from random numbers for the wave functions on the coarsest level. The accuracy of the wave functions is controlled by calculating the norms of the residuals of the eigenstates and it is finally improved by adding more V cycles starting from the finest level. A certain accuracy is needed in order to initiate self-consistency iterations that converge without large density oscillations. Then the new electron density and the ensuing effective potential are calculated. The new potential is not directly fed into the next iteration but it is mixed in this place, as well as later between the self-consistency iterations, with the input potential of the iteration. We monitor the accuracy of the wave functions by calculating their residuals and require that the accuracy has improved from the previous iteration. Usually one V cycle is sufficient for this, because the changes in the potential are small.

An important point is to also find a proper balance with respect to the pre- and postsmoothing sweeps on the different grid levels. Typically, on the finest level, we made two pre- and postsmoothing sweeps, whereas on the coarser grids their number is four. Actually, this means that on the

finest level, four successive sweeps are done if the potential is not updated. A potential update is always preceded and followed by two immediate smoothing sweeps.

IV. TESTS

We test the performance of our scheme by calculating the self-consistent electronic structures of a CO_2 molecule as well as that of perfect bulk Si lattice with a supercell of 64 Si atoms. The former system is an example of the employment of Dirichlet boundary conditions and the use of “hard” pseudopotentials, whereas the latter system represents the use of periodic boundary conditions and a supercell size typical in electronic structure calculations for point defects in solids.

The ions are described by pseudopotentials of the Kleinman-Bylander form,²⁷

$$V_{ion}(\mathbf{r}) = \sum_a V_{ion,loc}(|\mathbf{r}_a|) + \sum_{a,n,l,m} \frac{1}{\langle \Delta V_{lm}^a \rangle} |\Delta V_{ion,l}(r_a) u_{lm}(\mathbf{r}_a)\rangle \times \langle \Delta V_{ion,l}(r'_a) u_{lm}(\mathbf{r}'_a) |, \quad (23)$$

where $\langle \Delta V_{lm}^a \rangle$ is a normalization factor,

$$\langle \Delta V_{lm}^a \rangle = \int u_{lm}(\mathbf{r}_a) \Delta V_{ion,l}(r_a) u_{lm}(\mathbf{r}_a) d^3r, \quad (24)$$

and $\mathbf{r}_a = \mathbf{r} - \mathbf{R}_a$, u_{lm} are the atomic pseudopotential wave functions of angular and azimuthal momentum quantum numbers (l, m) , from which the l -dependent ionic pseudopotentials $V_{ion,l}(r)$ are generated using the Troullier-Martins scheme.²⁸ The ion core is assumed to be spherically symmetric. $\Delta V_{ion,l}(r) = V_{ion,l}(r) - V_{ion,loc}(r)$ is the difference between the l component of the ionic pseudopotential and the local ionic potential. We have chosen the s component of the pseudopotential as the local component.

Because the functions $|\Delta V_{ion,l}(r_a) u_{lm}(\mathbf{r}_a)\rangle$ are short ranged, operating on the wave function by the nonlocal parts of the pseudopotential is in practice a multiplication by a sparse matrix. The numerical work required to compute this scales as the square of the number of atoms in the system, whereas in the conventional reciprocal-space formulation, the work scales as the cube of the system size. The advantage of implementing the nonlocal pseudopotentials in real space has been noted also in the context of plane-wave methods.²⁹

In the previous multigrid implementations of the pseudopotential method,^{17,14} the nonlocal parts have only been employed on the finest grid. It is, however, straightforward to also implement them on the coarse levels, and we have found that this may increase the convergence rate and stability of the method.

The CO_2 molecule is placed diagonally in the center of a cubic computation volume of the size of $(12.6 a_0)^3$. The experimental C–O bond length of $2.19 a_0$ is used. Dirichlet boundary conditions are used so that the potential values outside the cube are obtained from a multipole expansion of




FIG. 2. Convergence of the total energy for the CO_2 molecule using direct mixing with different values of the feedback parameter κ ; $\kappa=0.4$ (solid line), $\kappa=0.5$ (dashed line), $\kappa=0.6$ (dash dotted line), and $\kappa=0.7$ (dotted line). A horizontal line has been added to indicate the chemical accuracy of 1 meV.

the charge density. The point mesh used is 63^3 , giving the grid spacing $h=0.20a_0$. The Mehrstellen discretization by Briggs *et al.*¹⁴ is used.

In this calculation we used a mixing scheme, where the new effective potential V_{in}^{i+1} is obtained from the input and output potentials according to

$$V_{in}^{i+1} = (1 - \kappa)V_{in}^i + \kappa V_{out}^i. \quad (25)$$

The convergence of the self-consistency iterations employing the strategy described above (Fig. 1) is shown in Fig. 2. The deviation of the total energy from the converged value is given as a function of self-consistency iteration steps performed. The zeroth iteration is a full-multigrid solution for the wave functions in the initial potential. Two V cycles starting from the finest level are included in this step. The effective potential obtained from the output electron density is mixed with the initial potential using the feedback $\kappa=0.4$. Next, at iteration one, the wave functions are relaxed in this new potential using one V cycle. From this point on, the four curves in the figure give the convergence with different values of the feedback parameter κ . One V cycle per self-consistency iteration step is done. A wide range of values for κ gives satisfactory convergence indicating a robust behavior for the scheme. The accuracy of 1 meV, which is sufficient in practical calculations, is reached after three or four V cycles. The implementation of the nonlocal parts of the pseudopotential on the coarse levels is found to speed up the convergence especially in this region. From Fig. 2 we obtain an average convergence rate of approximately one decade per self-consistency iteration. This is of the same order as those reported by Wang and Beck¹⁸ in their FAS scheme or by Kresse and Furthmüller²⁵ in their plane-wave




FIG. 3. Valence electron density in the (110) plane obtained in the Γ -point calculation for the 64-atom supercell of bulk Si. The area of the figure corresponds to the extent of the supercell.

scheme employing self-consistency iterations. The convergence rate of one decade per self-consistency iteration is better than that obtained by Ancilotto *et al.*¹⁷ in the FMG scheme and much better than the rate reached in the linearized multigrid scheme by Briggs *et al.*¹⁴

In our calculations for the CO_2 molecule, the initial full-multigrid step takes about three cpu minutes on a 667 MHz Compaq EV67 alpha processor. The succeeding V cycles take about one minute each. In order to put these figures into a proper context, we have solved the CO_2 molecule also by the plane-wave code CASTEP (Ref. 30) using a Fourier grid of 64^3 points. In both calculations, the errors in the total energy reduce at similar rates as a function of the cpu time.

We have solved for the electronic structure of perfect Si lattice described by a supercell of 64 Si ions. The lattice constant of $20.38 a_0$ used is the equilibrium value obtained in a plane-wave calculation, with which we have compared our real-space results. The first Brillouin zone is sampled in this test using the Γ point only. The point mesh used for the wave functions is 64^3 , giving the grid spacing $h=0.32a_0$. For the densities and potentials we use a finer grid of 128^3 points. The other numerical parameters and the iteration strategy are the same as in the CO_2 test. The resulting valence electron density on the (110) plane is given in Fig. 3. The area of the figure corresponds to the extent of the supercell. One notes that exactly the same features are reproduced at the equivalent points in different regions of the supercell. This means that a fully converged result has been found. We have compared the results of our real-space code to those obtained using the plane-wave method. The energy cutoff, 18 Ry, of the plane-wave expansion was chosen so that it results in a real-space point mesh of 64^3 , i.e., it is the same as in our real-space calculation. The widths of the valence band and band gaps obtained by the two methods agree with an accuracy of 3 meV. In the case of degenerate eigenstates, the real-space code results in degenerate eigenenergies with an accuracy better than 1 meV. The convergence towards the self-consistent solution occurs similarly as for the CO_2 mol-

ecule in Fig. 2. Thus, the convergence process seems to be independent of the size of the system.

V. SUMMARY AND OUTLOOK

In this paper we have generalized the RQMG method introduced by Mandel and McCormick²⁴ for the simultaneous solution of a desired number of lowest eigenenergy states. This approach can be viewed as belonging to a third group of multigrid methods, in addition to FAS and the techniques where the eigenproblem is linearized. In principle, one can use arbitrarily coarse grids in RQMG, whereas in the other multigrid methods, one has to be able to represent all the states on the coarsest grid.

We have demonstrated the feasibility of the method by electronic structure calculations for the CO₂ molecule and bulk Si described by pseudopotentials. Our strategy for the self-consistent solution consists of a full-multigrid solution for the wave functions in the initial potential, and subsequent self-consistency iterations. Less than five V cycles are generally sufficient for practically sufficient accuracy. The cpu times required for the FMG and SCF steps are roughly equal.

We have applied the method also in two-dimensional problems for quantum dots employing the current spin-density functional theory,³¹ in three-dimensional cylindrically symmetric systems, and also for calculation of positron states in solids.²⁶

We believe that our method will eventually compete with the standard plane-wave methods for electronic structure cal-

culations. However, some straightforward programming is still required. For calculations, where the optimization of the ionic structure is necessary, the Hellmann-Feynman forces will be implemented. In order to remove the spurious dependence of the total energy on the position of the atoms with respect to the grid points, Fourier filtering of the pseudopotentials is required.²⁹ Complex wave functions for any \mathbf{k} point are easily implemented, and are already in use in two-dimensional geometries.

Parallelization over \mathbf{k} points can be done easily. One only needs to communicate the electron density and effective potential at the end of each V cycle. During the RQMG V cycle, the states are all relaxed simultaneously and independently of each other. Therefore, parallelization over states is natural and easy to implement. However, for larger systems the Gram-Schmidt orthogonalization becomes very inefficient in a state-parallel code. The most efficient and yet straightforward choice is real-space domain decomposition,¹⁴ where each processor is mapped to a specific region of space.

ACKNOWLEDGMENTS

M.H. acknowledges stimulating discussions with Professor S. F. McCormick and Professor T. L. Beck. T.T. acknowledges financial support by The Vilho, Yrjö, and Kalle Väisälä Foundation. This research has been supported by the Academy of Finland through its Center of Excellence Program (2000–2005).

-
- ¹For a review, see R. O. Jones and O. Gunnarsson, *Rev. Mod. Phys.* **61**, 689 (1989).
- ²T. L. Beck, *Rev. Mod. Phys.* **72**, 1041 (2000).
- ³T. A. Arias, *Rev. Mod. Phys.* **71**, 267 (1999).
- ⁴M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias, and J. D. Joannopoulos, *Rev. Mod. Phys.* **64**, 1045 (1992).
- ⁵G. Kresse and J. Furthmüller, *Phys. Rev. B* **54**, 11 169 (1996).
- ⁶D. Bai and A. Brandt, *SIAM J. Sci. Comput. (USA)* **8**, 109 (1987).
- ⁷E. J. Bylaska, S. R. Kohn, S. B. Baden, A. Edelman, R. Kawai, M. E. G. Ong, and J. H. Weare, in *Proceedings of the 7th SIAM Conference on Parallel Processing for Scientific Computing*, edited by D. H. Bailey *et al.* (SIAM, Philadelphia, 1995), p. 219.
- ⁸J.-L. Fattebert, *J. Comput. Phys.* **149**, 75 (1999).
- ⁹F. Gygi and G. Galli, *Phys. Rev. B* **52**, 2229 (1995).
- ¹⁰N. A. Modine, G. Zumbach, and E. Kaxiras, *Phys. Rev. B* **55**, 10 289 (1997).
- ¹¹U. V. Waghmare, H. Kim, I. J. Park, M. Modine, P. Maragakis, and E. Kaxiras, cond-mat/0006183 (unpublished).
- ¹²S. Goedecker, *Rev. Mod. Phys.* **71**, 1085 (1999).
- ¹³J.-L. Fattebert and J. Bernholc, *Phys. Rev. B* **62**, 1713 (2000).
- ¹⁴E. L. Briggs, D. J. Sullivan, and J. Bernholc, *Phys. Rev. B* **54**, 14 362 (1996).
- ¹⁵A. Brandt, *Math. Comput.* **31**, 333 (1977).
- ¹⁶E. L. Briggs, D. J. Sullivan, and J. Bernholc, *Phys. Rev. B* **52**, 5471 (1995).
- ¹⁷F. Ancilotto, P. Blandin, and F. Toigo, *Phys. Rev. B* **59**, 7868 (1999).
- ¹⁸J. Wang and T. L. Beck, *J. Chem. Phys.* **112**, 9223 (2000).
- ¹⁹F. F. Grinstein, H. Rabitz, and A. Askar, *J. Comput. Phys.* **49**, 4423 (1983).
- ²⁰Y. Saad and M. H. Schultz, *SIAM J. Sci. Comput. (USA)* **7**, 856 (1986).
- ²¹A. Brandt, S. F. McCormick, and J. W. Ruge, *SIAM J. Sci. Comput. (USA)* **4**, 244 (1983).
- ²²T. L. Beck, K. A. Iyer, and M. P. Merrick, *Int. J. Quantum Chem.* **61**, 341 (1997).
- ²³S. Costiner and S. Ta'asan, *Phys. Rev. E* **52**, 1181 (1995).
- ²⁴J. Mandel and S. McCormick, *J. Comput. Phys.* **80**, 442 (1989).
- ²⁵G. Kresse and J. Furthmüller, *Comput. Mater. Sci.* **6**, 15 (1996).
- ²⁶A. P. Seitsonen, M. J. Puska, and R. M. Nieminen, *Phys. Rev. B* **51**, 14 057 (1995).
- ²⁷L. Kleinman and D. M. Bylander, *Phys. Rev. Lett.* **48**, 1425 (1982).
- ²⁸N. Troullier and J. L. Martins, *Phys. Rev. B* **43**, 1993 (1991).
- ²⁹R. D. King-Smith, M. C. Payne, and J. S. Lin, *Phys. Rev. B* **44**, 13 063 (1991).
- ³⁰M. C. Payne, M. P. Teter, D. C. Alan, T. A. Arias, and J. D. Joannopoulos, *Rev. Mod. Phys.* **64**, 1045 (1992). CASTEP is available from Molecular Simulations.
- ³¹G. Vignale and M. Rasolt, *Phys. Rev. B* **37**, 10 685 (1988).