

# AUSTRIAN GRID

Materialsimulationen auf atomarer Skala

Document Identifier:	<b>AG-DA-WP15-1-2009.pdf</b>
Status:	<b>Public</b>
Workpackage:	<b>15</b>
Partners:	
Lead Partner:	
WP Leaders:	<b>K.Schwarz, P.Blaha, (Inst. Materialchemie, TU Wien)</b>

### Delivery Slip

	Name	Partner	Date	Signature
From				
Verified by				
Approved by				

### Document Log

Version	Date	Summery of changes	Author

# MATERIALSIMULATIONEN AUF ATOMARER SKALA

P.Blaha

Institut  
Materialchemie, TU Wien  
pblaha@theochem.tuwien.ac.at

June 19, 2009

# Iterative Diagonalization in APW Based Methods in Electronic Structure Calculations

P. Blaha <sup>a,\*</sup>, H. Hofstätter <sup>b</sup>, O. Koch <sup>c</sup>, R. Laskowski <sup>a</sup>,  
K. Schwarz <sup>a</sup>

<sup>a</sup>*Institute of Materials Chemistry, Vienna University of Technology,  
Getreidemarkt 9/165-TC, A-1060 Vienna, Austria*

<sup>b</sup>*Institute of Mathematics, University of Vienna, Nordbergstrasse 8-10, A-1090  
Vienna, Austria*

<sup>c</sup>*Institute for Analysis and Scientific Computing, Vienna University of  
Technology, Wiedner Hauptstrasse 8-10, A-1040 Vienna, Austria*

---

## Abstract

Due to the increased computer power and advanced algorithms quantum mechanical calculations based on Density Functional Theory are more and more widely used to solve real materials science problems. In this context large nonlinear generalized eigenvalue problems must be solved repeatedly to calculate the electronic ground state of a solid or molecule. Due to the nonlinear nature of this problem, an iterative solution of the eigenvalue problem can be more efficient provided it does not disturb the convergence of the self-consistent-field problem. The blocked Davidson method is one of the widely used and efficient schemes for that purpose, but its performance depends critically on the preconditioning, i.e. the procedure to improve the search space for an accurate solution. For more diagonally dominated problems, which appear typically for plane wave based pseudopotential calculations, the inverse of the diagonal of  $(H - ES)$  is used. However, for the more efficient “augmented plane wave + local-orbitals” basis set this preconditioning is not sufficient due to large off-diagonal terms caused by the local orbitals. We propose a new preconditioner based on the inverse of  $H$  and demonstrate its efficiency for real applications using both, a sequential and a parallel implementation of this algorithm into our WIEN2k code.

*Key words:* Iterative diagonalization, Davidson method, Computational materials science, WIEN2k, APW, augmented plane waves  
*1991 MSC:* 65F10, 65F15, 65Z05, 81-08

---

\* Corresponding author.

*Email address:* pblaha@theochem.tuwien.ac.at (P. Blaha).

## 1 Introduction

In the last years quantum mechanical calculations based on Density Functional Theory [8] became very powerful due to the development of faster computers and more efficient algorithms. The key problem of such calculations is the solution of the Kohn-Sham (KS) equation [12]

$$\hat{H}\Psi_i(r) = \left[ -\frac{1}{2}\nabla^2 + v_{eff}(r) \right] \Psi_i(r) = \epsilon_i\Psi_i(r) \quad (1)$$

where the Hamilton operator  $\hat{H}$  contains the kinetic energy operator and an effective potential  $v_{eff}(r)$ . The latter depends on the electron density  $\rho(r)$ , which can be obtained from the wave functions  $\Psi_i$ , but obviously  $\Psi_i$  itself depends on  $v_{eff}(r)$  and thus on  $\rho(r)$ . Hence, an iterative solution of the KS equations must be found by fixed point iteration, i.e. starting with some initial guess for  $\rho(r)$  and iterating until the input and output density,  $\rho_{in}(r)$  and  $\rho_{out}(r)$ , becomes nearly the same (within a threshold). This is called the self-consistent-field (scf) procedure.

Each wave function  $\Psi$  is usually expanded in a set of (non-orthogonal) basis functions  $\phi(r)$

$$\Psi(r) = \sum_j c_j \phi_j(r) \quad (2)$$

and the best wave functions are determined by a Rayleigh-Ritz procedure using the variational principle which leads to a generalized eigenvalue problem (see Section 2, Eq. 3). The dimension of this problem depends on the investigated system (mainly on the number of atoms) and on the choice of the basis set  $\phi$ , but typically varies for large scale problems from  $10^3$  up to  $10^5$ . As mentioned above, the KS equations have to be solved iteratively, thus the corresponding eigenvalue problem does not need to be solved exactly at the beginning of the scf-iterations, but an approximate solution is fine as long as it does not affect the convergence of the scf procedure. Note that the changes in the eigenvalue problem become smaller and smaller once the scf-procedure is nearly converged since  $v_{eff}$  will hardly change anymore. In addition, for most applications not all eigenvalues are needed, but for instance in the linearized augmented plane wave (LAPW), or its more efficient variant, the augmented plane wave plus local orbital (APW+lo) method implemented in the WIEN2k code [3,17,2,19,20], only the lowest 5-20 % of all  $\epsilon_i$  are needed.

The problem as sketched above calls for an iterative diagonalization scheme since an accurate solution of the eigenvalue problem is only necessary at the very end of the scf-procedure and thus significant computer time can

be saved in comparison to standard diagonalization procedures (this will be called “full diagonalization” later in this paper) based on LAPACK routines [9] for Cholesky factorization of the overlap matrix  $S$  (zpotrf), reduction to standard eigenvalue problem (zhegst), tri-diagonalization and solving for the lowest  $k$  eigenvalues and eigenvectors (zheevx) and finally backtransformation of the eigenvectors to the original generalized eigenvalue problem (ztrsm).

The iterative method for the diagonalization of generalized eigenproblems implemented so far in the WIEN2k code [3] is a blocked version of the Davidson method [4,5] which was introduced in [21]. Iterative methods for the problem at hand are also discussed in [1,13–15,18,24,26]:

In [26], the method of *RMM-DIIS* (*residual minimization/direct inversion in the iterative subspace*) is proposed and compared with the *Davidson* and *Block Davidson* methods. The latter has the disadvantage that the doubling in the dimension of the search space is prohibitive for large initial subspaces. Therefore the RMM-DIIS method is claimed to have the advantage that only matrices of the size of the number of previous iteration steps are necessary. However, the method is fundamentally sequential in nature which the authors recognize as a major drawback, and which in the light of the development of parallel and grid-enabled versions of the WIEN2k code makes this approximate diagonalization unattractive. A comparison with several other methods shows that (disregarding computational cost) the block Davidson method displays the best improvement in accuracy per iteration step due to the doubling of the search space [26].

Ref. [24] gives an overview of the state of the art of iterative diagonalization at that time, and demonstrates that a new *preconditioned conjugate gradient method* compares most favorably with *conjugate gradients*, *steepest descent* and *imaginary time propagation*. The VASP code [15], which is a highly efficient plane wave pseudopotential code, uses the RMM-DIIS method of [26] in a variant proposed in [18]. They claim this method to be superior for very large problems [13] if an unblocked, band-by-band iteration is used.

In this paper, we propose an approximate diagonalization which is motivated by the fact that the Davidson method previously implemented in the WIEN2k code [21] has proven to be no longer satisfactory once the basis set has been changed from the standard LAPW to the APW+lo basis set [17]. Apparently, the underlying discretization, the importance of non-diagonal terms (the local orbital contribution to the plane wave basis) and the adaptive basis set (the basis set changes slightly with  $v_{eff}$ ) makes the preconditioning with only the diagonal elements  $diag^{-1}[H - ES]$  not efficient anymore. Our new method is motivated by the improvements promised by the Jacobi-Davidson method [22,23,25] as compared to the original Davidson method [5]. However, application of the subspace expansion from the Jacobi-Davidson method seems

prohibitively expensive, hence we propose a simplification which uses an approximate computation of the subspace expansion related to an iterative solution of the associated linear system of equations. Furthermore, we are going to demonstrate in Section 4 that the method is superior to full diagonalization in both efficiency and parallelism.

## 2 New Diagonalization Method

We want to compute the  $k$  eigenpairs corresponding to the lowest eigenvalues of the generalized eigenproblem

$$HY = SY\Lambda, \quad (3)$$

where  $H \in \mathbb{C}^{n \times n}$  is hermitian and  $S \in \mathbb{C}^{n \times n}$  is hermitian and positive definite<sup>1</sup>. The solution of (3) consists of the (matrix of the) lowest eigenvalues  $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_k)$  and matrix  $Y \in \mathbb{C}^{n \times k}$  containing  $k$  corresponding eigenvectors  $y_j$ . The algorithm proceeds as follows:

- Input:

$$Y = [y_1, \dots, y_k] \in \mathbb{C}^{n \times k}. \quad (4)$$

Usually these are approximations to eigenvectors which were computed in the last SCF cycle.

- Compute the Ritz values (Rayleigh quotients)

$$\vartheta_j = \frac{y_j^\dagger H y_j}{y_j^\dagger S y_j}, \quad j = 1, \dots, k. \quad (5)$$

- Set up the search space  $[Y \ Z] \in \mathbb{C}^{n \times 2k}$  with

$$z_j = H^{-1}(H - \vartheta_j S)y_j, \quad j = 1, \dots, k. \quad (6)$$

- Set up the reduced problem

$$\tilde{H}V = \tilde{S}V\Gamma, \quad (7)$$

where

$$\tilde{H} = [Y \ Z]^\dagger H [Y \ Z] = \begin{bmatrix} Y^\dagger H Y & Y^\dagger H Z \\ Z^\dagger H Y & Z^\dagger H Z \end{bmatrix} \quad (8)$$

<sup>1</sup> This implies that there exists a basis of eigenvectors to the real eigenvalues  $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$  which are orthonormal w. r. t. the inner product defined by  $S$ .

and

$$\tilde{S} = [Y \ Z]^\dagger S [Y \ Z] = \begin{bmatrix} Y^\dagger S Y & Y^\dagger S Z \\ Z^\dagger S Y & Z^\dagger S Z \end{bmatrix}. \quad (9)$$

- Compute eigenvectors  $V_{1:k}$  of (7) corresponding to the lowest  $k$  eigenvalues  $\gamma_1 \leq \dots \leq \gamma_k$ . Here and in the following,  $V_{1:k}$  refers to the first  $k$  columns of the matrix  $V \in \mathbb{C}^{n \times m}$ . We may assume that  $V_{1:k}$  is orthonormal with respect to  $\tilde{S}$ , i.e.,  $V_{1:k}^\dagger \tilde{S} V_{1:k} = I_k$ , where  $I_k$  denotes the identity matrix of dimension  $k$ .
- Compute new approximations

$$Y_{\text{new}} = [Y \ Z] V_{1:k} \quad (10)$$

to the eigenvectors of (3).

*Remarks:*

- In the original Davidson method [5], the subspace is expanded according to

$$z_j = \text{diag}^{-1}(H - \vartheta_j S)(H - \vartheta_j S)y_j, \quad j = 1, \dots, k,$$

while in the full Jacobi-Davidson method [25], the subspace is expanded according to

$$(I_n - S y_j y_j^\dagger)(H - \vartheta_j S)(I_n - y_j y_j^\dagger S)z_j = (H - \vartheta_j S)y_j. \quad (11)$$

The subspace expansion in (6) corresponds to the first step of an iterative solution of the linear equations system (11) and one can rationalize this choice if the entries of  $H$  are larger than those of  $S$  (which is always true in our cases).

- Note that in (6) inversion of  $H$  is possible only for nonsingular matrices. We have not met any problem with this restriction in practice, but in order to avoid vanishing eigenvalues, the generalized eigenvalue problem (3) can equivalently be reformulated by letting  $H \leftarrow H + \varepsilon S$ , for any  $\varepsilon \in \mathbb{R}$ . This shifts the eigenvalues,  $\lambda_j \leftarrow \lambda_j + \varepsilon$ , but leaves the eigenvectors unaltered. Thus, (3) can be assumed without restriction of generality to be stated with a nonsingular matrix  $H$ .
- If the SCF cycle is nearly converged, the initial guesses  $y_j$  may almost coincide with the corresponding exact eigenvectors of (3). In that case, the residual  $(H - \vartheta_j S)y_j$  is zero, and the reduced system matrices  $\tilde{H}$  and  $\tilde{S}$  are singular, hence the reduced problem (7) admits no unique solution. This causes problems with the LAPACK routines used for the “small” (reduced) eigenvalue problem (7) in the parallel version. Our work-around to refine the remaining eigen-pairs is to *project out* the converged eigenvalues and associated -vectors. Thus, if the index  $j$  corresponds to a converged state,

then we set the associated entries  $\tilde{S}_{j,i} = \tilde{H}_{j,i} := 0$ ,  $\tilde{S}_{i,j} = \tilde{H}_{i,j} := 0$ ,  $i = 1, \dots, m$ ,  $\tilde{H}_{j,j} = 10^8$ ,  $\tilde{S}_{j,j} = 1$ . This moves the eigenvalue associated with the index  $j$  beyond the physically relevant range and thus this contribution no longer appears in the further computations.

### 3 Implementation Details

The implementation of the iterative diagonalization scheme described above (Eq. 3-10) is based on highly optimized code using BLAS and LAPACK (or SCALAPACK) routines. It can solve either a real symmetric or a complex hermitian eigenvalue problem (all in double precision) in sequential or (mpi-) parallel mode.

In order to set up the reduced eigenvalue problem (7) we first form the products of  $HY$  and  $SY$  (zhemm), then evaluate the upper part of the reduced Hamiltonian and overlap matrix  $Y^\dagger HY$ ,  $Y^\dagger SY$  (zher2k). The diagonals of these matrices are used to compute the Ritz values  $\vartheta_j$  (5). The search space represented by the matrix  $Z$  defined in (6) is a solution of the linear system of equations:

$$Hz_j = (H - \vartheta_j S)y_j, \quad j = 1, \dots, k \quad (12)$$

which is solved (zhetrs) after factorization of  $H$  (zhetrf). The products of  $HY$  and  $SY$  that enter the right-hand side of this equation have been obtained before, and the right-hand side itself can be used to evaluate (8) and (9),  $Z^\dagger HY$ ,  $Z^\dagger SY$ ,  $Z^\dagger HZ$  and  $Z^\dagger SZ$  (zgemm). The reduced eigenproblem (7), which is only of dimension  $2 \cdot k$ , is usually solved quickly by full diagonalization as described in the introduction.

### 4 Numerical Results

We start this section by giving a decomposition of the total CPU time of the iterative scheme into the different tasks. As shown in Fig.1 most of the time is spent on the factorization of the  $H$  matrix (the LAPACK solver for linear equations requests factorized  $H$ ) and in the actual solution of the linear equation system (12). These tasks scale as  $N^3$  and  $N^2 \times k$ , but the prefactor of factorization is much smaller than the one for the full diagonalization [7]. A considerable fraction of the CPU time, especially for small matrices, is also spent on the multiplication of the  $H$  or  $S$  matrices with the  $Y$  or  $Z$  vectors, which also scales as  $N^2 \times k$ . For the current algorithm we need three such

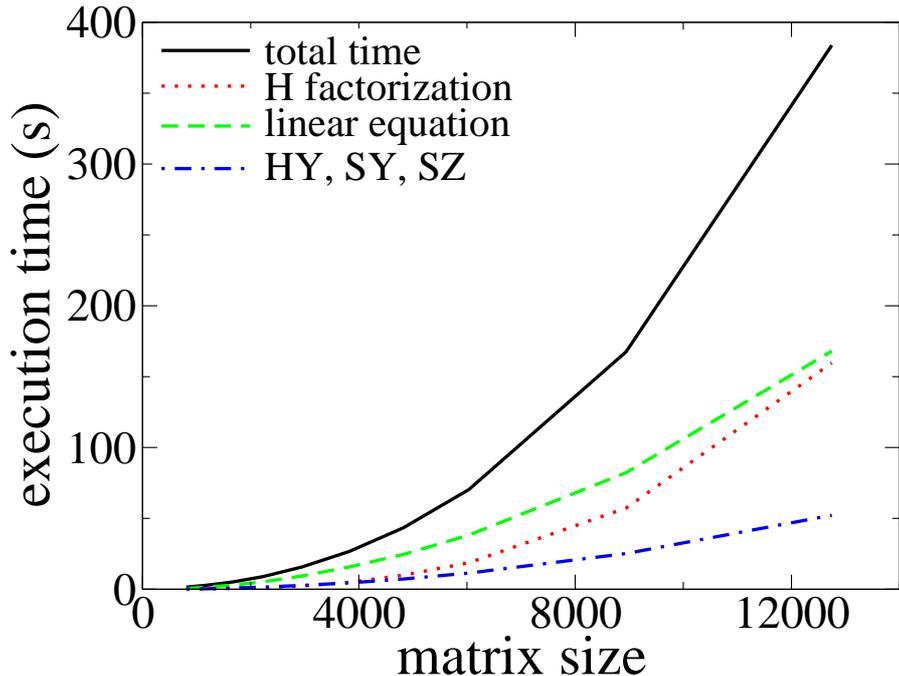


Fig. 1. CPU time analysis of our iterative diagonalization scheme of a real symmetric eigenvalue problem as function of matrix size and 298 eigenvalues (see text)

multiplications:  $HY$ ,  $SY$  and  $SZ$ . (Note that the product  $HZ$  is a combination of  $HY$  and  $SY$ ). The CPU time spent on those three products is also indicated in Fig. 1. All other parts, including the solution of the reduced eigenvalue problem (7) are usually much smaller, because  $(2 \cdot k)^3$  is never large for realistic cases.

As discussed above, the CPU time of our iterative algorithm scales proportionally to  $\alpha N^3 + \beta(N^2 \times k)$ , where  $N$  is the size of  $H$  and  $k$  is the number of evaluated eigenvalues. Since the full diagonalization scales proportionally to  $\alpha' N^3$  the gain in the CPU time depends on the ratio of  $\alpha$  and  $\alpha'$  and on how many eigenvalues are needed. The actual speedup is indicated in Fig. 2, where the ratio of the CPU time of full and iterative diagonalization is shown for 3 different values of the number of eigenvalues  $k$  (and as a function of the matrix size  $N$ ).

As mentioned in the introduction, our ultimate goal is to solve the KS equations for a particular solid or molecule self-consistently. Thus, to evaluate the efficiency and usefulness of this iterative diagonalization one should in fact consider the gain in time for one discrete scf-cycle, i.e. not only the time for diagonalization but also for setting up the eigenvalue problem as well as other parts like the calculation of the electron density  $\rho$  and the potential  $v_{eff}$ . Furthermore, one has to verify that the iterative diagonalization scheme does not increase the number of scf-cycles that are necessary to converge  $\rho$ . We will

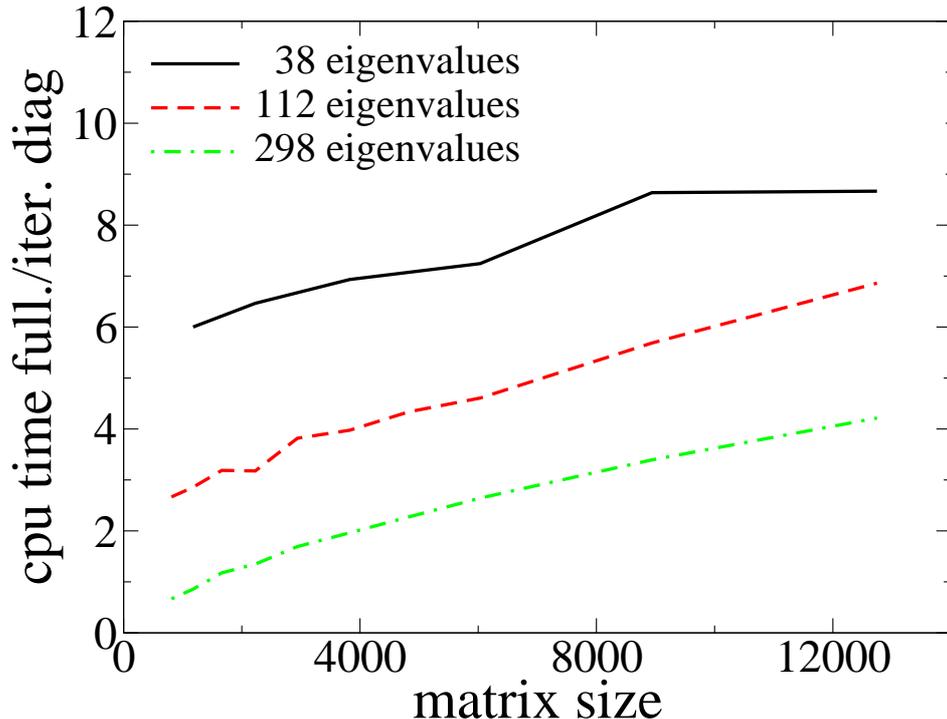


Fig. 2. Ratio of the CPU time of full and iterative diagonalization as function of matrix size for 3 different numbers of eigenvalues

demonstrate this on two very different examples below:

Since WIEN2k is a code utilizing a finite unit cell with 3-dimensional periodic boundary conditions, a calculation of a single molecule must utilize large cells so that the molecules are well separated from each other. This leads to a rather large number of APW basis functions. We have done such calculations for the CO molecule and such an example can be considered as extreme case where, on the one hand, the time for setting up the eigenvalue problem ( $H, S$ ) is rather negligible compared to the diagonalization, because they scale like  $N_{bas}^2 * N_{atoms}$ , where the number of basis functions  $N_{bas}$  is 2800 while  $N_{atoms}$  is just 2 for this example. The eigenvalue problem is complex hermitian, because there is no inversion symmetry present. On the other hand only one k-point in the Brillouin zone (BZ) needs to be considered, i.e. only one eigenvalue problem needs to be solved in every scf cycle and thus there is a relatively large fraction of time spent in the calculation of the potential  $v_{eff}$ . Furthermore, the number of required eigenvalues  $\epsilon_i$  will be extremely small (below 1 % of the matrix size). In fact as can be seen from Table 1 about 30 % of the time goes into the calculations of  $\rho$  and  $v_{eff}$  but 65 % into the diagonalization. Using our new iterative scheme, the diagonalization time can be reduced by a factor of 10, so that the eigenvalue problem becomes only a small fraction of the total time of one scf-cycle. As noted above, it is important that an iterative diagonalization scheme does not increase the number of scf-cycles too much, otherwise the gain in speed may be lost in the additional cycles required to

task	time (s)	iterations	time scf (s)
$\rho, v_{eff}$	24.7	-	-
$H, S$	2.3	-	-
full diag.	49.5	10	765.0
iter. diag (8 $\epsilon_i$ )	4.5	24	756.0
iter. diag (16 $\epsilon_i$ )	4.8	14	445.2
iter. diag (32 $\epsilon_i$ )	5.4	11	356.4

Table 1

Timings of different parts (see text) of the calculations for a CO molecule. The scf-timings take into account the required number of iterations to reach a convergence of  $5 \cdot 10^{-5}$  for the charge distance.

reach self-consistency. In Fig. 3 we show the charge distance (the integral  $|\rho_{in} - \rho_{out}|$ ) as function of the number of scf-iterations. The CO molecule has 5 occupied eigenstates and when we calculate just 8 states (i.e. only 3 unoccupied states), the number of cycles would increase drastically and there would be only a very small gain in time. On the other hand, with 16 or 32 eigenvalues (this does hardly affect the iterative diagonalization time) the required number of scf-cycles does not increase significantly and a large saving of total cpu-time can be obtained (see Table 1).

As a representative example of a typical medium-sized bulk calculation, we next present results for Stibnite,  $Sb_2S_3$ , a compound with 20 atoms/cell and inversion symmetry, thus the matrix elements are real numbers with  $N_{bas} = 2400$ . Typically one has to use about 10 k-points in the BZ, i.e. we have to solve 10 different eigenvalue problems in every scf-cycle, which has been taken into account for the corresponding scf timings listed in Table 2. Setting up the eigenvalue problem is here a significant time (37 % for full, but even up to 72 % for iterative diagonalization) and in fact even dominates when the iterative diagonalization is used. The gain in speed is not so large, because of the relatively small matrix and the fair amount of eigenvalues (about 10 %) required in this case, but still the total cpu time of the scf-procedure can be reduced by more than 25 %. The convergence of the scf-cycles depends only mildly on the number of eigenvalues and a number typically twice the number of occupied bands (96 in this case) is more than sufficient (see Fig. 4).

For large scale calculations (more than 64 atoms / unit cell) the dimension of the eigenvalue problem becomes so large that it will either take very long time on a single processor or eventually the matrices cannot even be stored in memory on a single node. Thus, fine grained, mpi-based parallelization distributing the matrices over many processors and using ScaLAPACK routines [10] for linear algebra operations is necessary. In the following we present results for a 3x3 super-cell of a h-BN/Ni(111) surface model, which contains

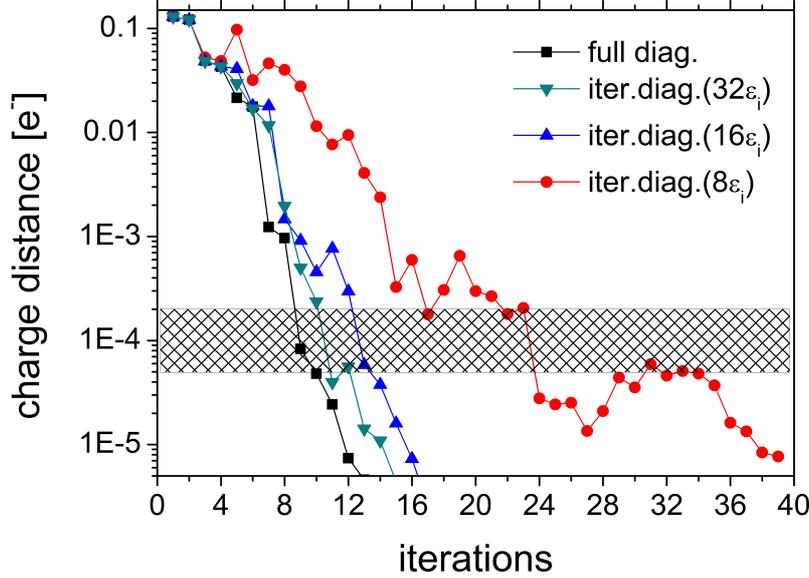


Fig. 3. Charge distance (integral  $|\rho_{in} - \rho_{out}|$ ) as function of the number of scf-iterations for the CO molecule using full or iterative diagonalization (with 8, 16 or 32 eigenvalues). The hatched region indicates the typically desired convergence criterion.

task	time (s)	iterations	time scf (s)
$\rho, v_{eff}$	16.0	-	-
$H, S$	5.2	-	-
full diag.	9.0	13	2054
iter. diag (150 $\epsilon_i$ )	2.1	17	1513
iter. diag (216 $\epsilon_i$ )	2.8	15	1440
iter. diag (361 $\epsilon_i$ )	5.8	13	1638

Table 2

Timings of different parts (see text) of the calculations for  $Sb_2S_3$ . The scf-timings take into account 10 k-points and the required number of iterations to reach a convergence of  $5 \cdot 10^{-5}$  for the charge distance.

99 atoms/cell. Since this is a magnetic and metallic system, 8 k-points and a spin-polarized calculation (i.e. 16 different eigenvalue problems per scf-cycle) had to be performed. The basis set size  $N_{bas} = 16900$  (real matrix elements) is so large, that we could not run this example on a single cpu, but the basic timings will be given for 4 cpus (cluster of AMD dual-core opteron with infiniband network). For the iterative diagonalization we used 1040 eigenvalues

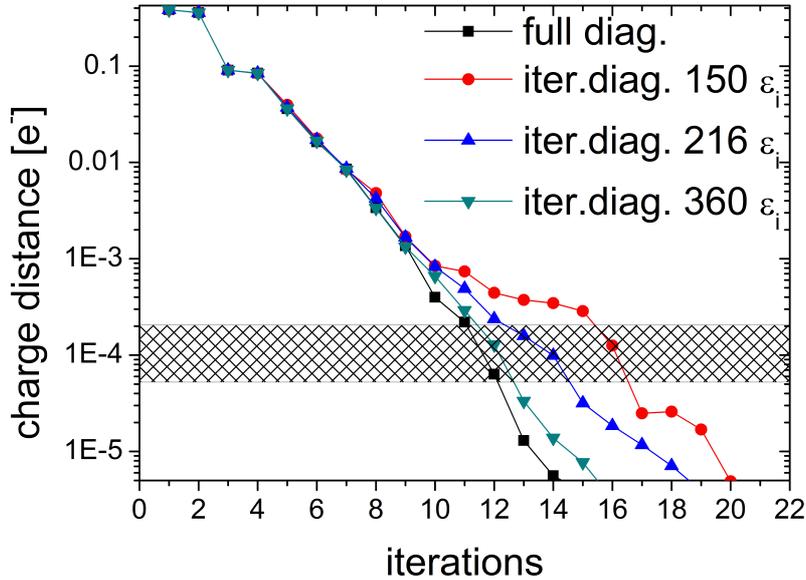


Fig. 4. Charge distance (integral  $|\rho_{in} - \rho_{out}|$ ) as function of the number of scf-iterations for  $Sb_2S_3$  using full or iterative diagonalization (with 150, 216 or 361 eigenvalues). The hatched region indicates the typically desired convergence criterion.

(about 600 are occupied) and virtually no scf-convergence degradation was found. As evident from Table 3, the time spent in calculating  $\rho$  and  $v_{eff}$  is almost negligible, but 70 % of the time is spent during full diagonalization. The iterative scheme is more than five times faster and thus only 30 % of the time is spent on the diagonalization. Overall the time of an scf cycle can be reduced by more than 55 %. Still the cpu time is so large that further parallelization is desirable and in Fig. 5 the scaling of setup and diagonalization of this eigenvalue problem up to 100 cpus is presented. It should be noted that these measurements were done under full-load conditions of the machine, which leads to some irregularities in particular when more than 64 cpus were used. The setup of the matrix elements is split into a spherical part  $H_{sph}, S_{sph}$ , which takes about 2/3 of the total setup time and a non-spherical part,  $H_{nsp}$  and  $S_{nsp}$ .  $H_{sph}, S_{sph}$  scales nearly perfectly as it has basically no communication and the sequential part takes only a fraction of a percent. The calculation of the non-spherical matrix elements  $H_{nsp}$  and  $S_{nsp}$  still scales very well, but as one cannot avoid some sequential steps, decreasing efficiency is to be expected for more cpus. The iterative diagonalization scales slightly better than full diagonalization, but is of course worse than the efficiency (speedup / number of processors) of the matrix setup. The efficiency of the full (iterative) diagonalization is 71 (82) % for 16 processors, 52 (62) % for 64 processors, and 40 (49) % for 100 processors. Note that for the latter the size of the sub-blocks

task	time (s)
$\rho, v_{eff}$	151
$H, S$	1176
full diag.	2454
iter. diag. (150 $\epsilon_i$ )	461
one scf cycle (full diag.)	57604
one scf cycle (iter. diag. 150 $\epsilon_i$ )	25616

Table 3

Timings on 4-processors of different parts (see text) of the calculations for a 3x3 super-cell of a h-BN/Ni(111) surface model with 99 atoms/cell ( $N_{bas} = 16900$ ). The scf-timings assume 16 k-points.

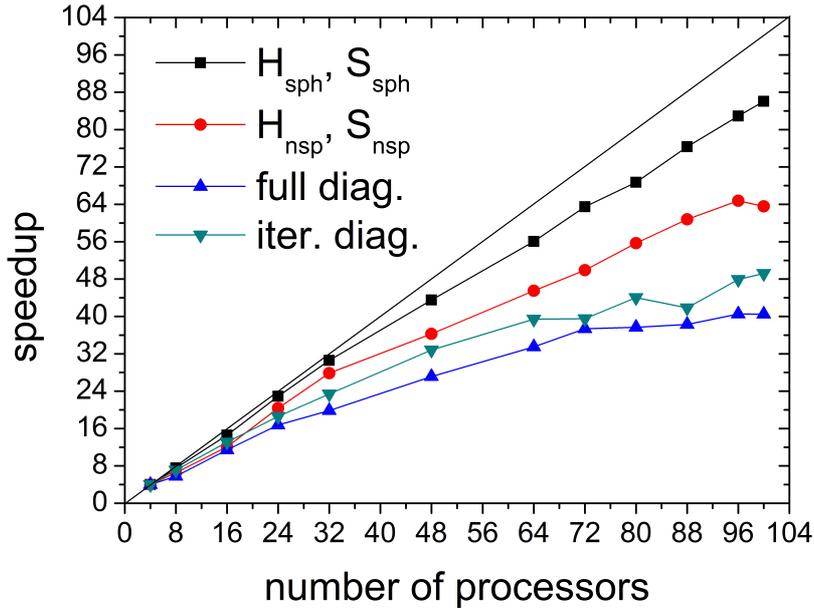


Fig. 5. Speedup of various parts of the eigenvalue problem with the number of processors for a 3x3 super-cell of a h-BN/Ni(111) surface model with 99 atoms/cell ( $N_{bas} = 16900$ ).

in the ScaLAPACK distribution is only 1700x1700 and thus comes close to the limit of efficient ScaLAPACK parallelizations.

## 5 Conclusions

When one needs to solve the KS equations self-consistently, one has to solve a sequence of closely related generalized eigenvalue problems thus an iterative solution could lead to a significant speed-up. We propose a new preconditioner for the blocked Davidson method to solve efficiently generalized eigenvalue problems also for the case of non-diagonally dominant matrices. Commonly the inverse of the diagonal of  $(H - ES)$  is used as preconditioner, but we use the inverse of  $H$  instead. The algorithm has been implemented in both, a sequential and a parallel version of our WIEN2k code and we demonstrate that it leads to a fast, stable and efficient scheme. In addition the new iterative scheme shows better scaling properties than the full diagonalization. Most importantly, the approximate solution of the eigenvalue problem does not affect the solution of the self-consistent-field problem (i.e. the number of scf iterations) provided a sufficient number of states is included.

This scheme has already been tested on an everyday basis by the large WIEN2k community [11] for several month and allowed us to increase the size of a materials science problem by a factor of two to three. Recently we could solve a real nanoscience problem with more than 1100 atoms/unit cell [6,16].

## Acknowledgements

This work was supported by the Austrian Grid project (WP 15).

## References

- [1] C. Bendtsen, O. Nielsen, and L. Hansen. Solving large nonlinear generalized eigenvalue problems from density functional theory calculations in parallel. *Appl. Numer. Math.*, 37:189–199, 2001.
- [2] P. Blaha, K. Schwarz, and G. Madsen. Electronic structure calculations of solids using the WIEN2k package for material sciences. *Comp. Phys. Commun.*, 147:71–76, 2002.
- [3] P. Blaha, K. Schwarz, G. Madsen, D. Kvasnicka, and J. Luitz. An augmented plane wave plus local orbital program for calculating crystal properties, 2001. ISBN 3-9501031-1-2.
- [4] M. Crouzeix, B. Philippe, and M. Sadkane. The Davidson method. *SIAM J. Sci. Comput.*, 15:62–76, 1994.

- [5] E. Davidson. The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices. *J. Comput. Phys.*, 17:87–94, 1975.
- [6] H. Dil, J. Lobo-Checa, R. Laskowski, P. Blaha, S. Berner, J. Osterwalder, and T. Greber. Surface trapping of atoms and molecules with dipole rings. *Science*, 319:1824, 2008.
- [7] G. H. Golub and C. F. Van Loan. *Matrix Computations*. The John Hopkins University Press, Baltimore, 2nd edition, 1989.
- [8] P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Phys. Rev.*, 136:B864, 1964.
- [9] <http://www.netlib.org/lapack/>.
- [10] <http://www.netlib.org/scalapack/>.
- [11] <http://www.wien2k.at>.
- [12] W. Kohn and L.J. Sham. Self-consistent equations including exchange and correlation effects. *Phys. Rev.*, 140:A1133, 1965.
- [13] G. Kresse and J. Furthmüller. Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set. *Comput. Mater. Sci.*, 6:15–50, 1996.
- [14] G. Kresse and J. Furthmüller. Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B*, 54:11169–11186, 1996.
- [15] G. Kresse and J. Furthmüller. VASP the GUIDE, 2007. Available at <http://cms.mpi.univie.ac.at/VASP>.
- [16] R. Laskowski and P. Blaha. Unraveling the structure of the h-bn/rh(111) nanomesh with ab initio calculations. *J. Phys.: Condens. Matter*, 20:064207, 2008.
- [17] G.K.H. Madsen, P. Blaha, K. Schwarz, E. Sjöstedt, and L. Nordström. Efficient linearization of the augmented plane-wave method. *Phys. Rev. B*, 64:195134, 2001.
- [18] P. Pulay. Convergence acceleration of iterative sequences. the case of scf iteration. *Chem. Phys. Lett.*, 73:393, 1980.
- [19] K. Schwarz. DFT calculations of solids with LAPW and WIEN2k. *Solid State Commun.*, 176:319–328, 2003.
- [20] K. Schwarz and P. Blaha. Solid state calculations using WIEN2k. *Comput. Mater. Sci.*, 28:259–273, 2003.
- [21] D. Singh. Simultaneous solution of diagonalization and self-consistency problems for transition-metal systems. *Phys. Rev. B*, 40:5428–5431, 1989.

- [22] Gerard L. G. Sleijpen, Albert G. L. Booten, Diederik R. Fokkema, and Henk A. Van der Vorst. Jacobi-Davidson type methods for generalized eigenproblems and polynomial eigenproblems. *BIT*, 36(3):595–633, 1996.
- [23] Gerard L. G. Sleijpen and Henk A. Van der Vorst. A Jacobi-Davidson iteration method for linear eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 17(2):401–425, 1996.
- [24] M. P. Teter, M. Payne, and D. C. Allan. Solution of Schrödinger’s equation for large systems. *Phys. Rev. B*, 40(18):12255, 1989.
- [25] H. J. J. van Dam, J. H. van Lenthe, G. L. G. Sleijpen, and H. A. van der Vorst. An improvement of Davidson’s iteration method. *J. Comput. Chem.*, 17, 3:267–272, 1996.
- [26] D.M. Wood and A. Zunger. A new method for diagonalising large matrices. *J. Phys. A: Math. Gen.*, 18:1343–1359, 1985.