Preconditioning of self-consistent-field cycles in density-functional theory: The extrapolar method

P.-M. Anglade and X. Gonze
Unité PCPM, Université Catholique de Louvain, Pl. Croix du Sud 1, 1348 Louvain-la-Neuve, Belgium
(Received 7 September 2007; revised manuscript received 2 May 2008; published 31 July 2008)

The number of self-consistent-field iterations needed for the density-functional theory treatment of metallic systems grows with the size of the unit cell, not only for basic algorithms like simple mixing, but also for more advanced schemes, in which results from several past steps are mixed. Preconditioning techniques have the potential to suppress this growth, although the available methods have strong limitations: either they deliver little improvement in case of mixed systems with metallic and nonmetallic regions, or the computation of the preconditioner scales badly with the size of the system, with a large prefactor. We propose an approximate preconditioner, with tremendously reduced prefactor, that makes the number of self-consistent cycle nearly independent of the size of the system, and bears little overhead up to the one hundred atom range. The susceptibility matrix, a key ingredient in our scheme, is approximated thanks to the closure relation. Instead of using the exact formulation of the dielectric matrix, we rely on the random-phase approximation, that allows to further decrease the prefactor thanks to a very low wave vector cutoff, even for systems with both vacuum and a metallic region. We test this algorithm for systems of increasing size and demonstrate its practical usefulness.

DOI: 10.1103/PhysRevB.78.045126 PACS number(s): 71.15.−m

I. INTRODUCTION

With the advent of both powerful computers and fast algorithms for solving the system of equations from density-functional theory DFT,[1] including so-called order (N) methods,[2] the spectrum of DFT application has expanded tremendously. While ten years ago, systems of 400 atoms seemed exceptionally big,[3] today, people consider systems with tens of thousand atoms or bigger.[4]

This idyllic picture of advances in DFT computation leaves some problems aside. In particular, since the very beginning of DFT, self-consistency has been an issue. Being able to run many times across the DFT Kohn and Sham equation is not enough if the process is not converging to the ground state. Even if the procedure converges, the increase of the number of self-consistency cycles needed to reach convergence might be overwhelming. Many papers (e.g., Refs. 5–10) propose methods to enforce or to speed up the DFT self-consistency convergence. However, none of them provides a completely satisfactory answer to the challenge of present and future large-scale DFT calculations for all classes of systems.

The simplest approach to self-consistency relies on a linear mixing of the output of an iteration with its input, used to start the next iteration. Although not the most efficient algorithm, it is particularly robust, as for each problem, there exists a sufficiently small mixing factor that will make the iterative procedure converge to a fixed point. However, in the case of systems with a metallic electronic spectrum (vanishing energy gap), such mixing factor has to be decreased with the size of the system, with a concomitant increase of the number of iterations needed to reach a sufficiently accurate solution.

In the present paper, we will focus on that class of systems, for which the number of iterations is roughly proportional to the square of the largest linear dimension of the system (L) within the simple mixing approach.[3] The problem is not that severe in the case of systems with an insulating electronic spectrum, as for such systems, the growth of the number of iterations saturates. Still, this number might be large, and ought to be decreased.

Actually, finding the ground state of the DFT equations is nothing else, mathematically speaking, than the optimization of a continuous functional in a large space of varying parameters. Starting from a guess close to the sought solution, the problem reduces to the optimization of a convex, nearly quadratic, functional, a well-known subject of mathematical analysis research. In order to improve upon simple mixing, two lines of thoughts might be followed, and even combined advantageously, that both arise from basic mathematical analysis of this problem: (1) taking into account the result of past iterations, and (2) using the available knowledge of the response of the system through a preconditioning scheme.

In the context of DFT, both lines of thoughts have been examined. Using the result of past iterations is the focus of Anderson,[5] Pulay[7] or conjugate-gradient,[11] and many other[12–14] schemes. Such schemes, that we will term "advanced mixing" algorithms, have theoretically the potential to bring the scaling of the number of iterations down from L^2 to L. On the other hand, preconditioning schemes, such as those proposed by Kerker[8] or Ho, Ihm, and Joannopoulos (HIJ, Ref. 9), and later developed by Sawamura *et al.*[15] and Auer *et al.*,[16,17] could theoretically suppress completely this scaling. While advanced mixing algorithms provide their improvement nearly irrespective of the type of system that is considered, this is not the case of preconditioning schemes. For existing preconditioning algorithms, either their range of application is limited (Kerker's scheme brings almost no improvement for strongly inhomogeneous systems), or the time to compute the preconditioner scales badly, and quickly becomes the dominating factor in the computation, as in HIJ technique.

Referring to N as a conveniently chosen variable describing the three-dimensional size of the system (for instance, the number of electrons, the cell volume or the number of atom—if the cell is cubic, N will be proportional to L^3), achieving overall order (N) DFT scaling implies order (1) for

the number of SCF cycles, because the computation of the density for a fixed potential cannot be reduced further than order (N) scaling. Moreover, the extra computational cost introduced by the use of the preconditioner should scale linearly at most. In the case of more standard order (N^3) algorithms, the constraint on the scaling of the extra computational cost due to the computation of the preconditioner might be weaker—it should also be an order (N^3) algorithm at most.

The HIJ method is based on the computation of the *exact* first-order preconditioner, the inverse of the dielectric matrix. As a consequence, when the starting point of the self-consistency is in the quadratic part of the functional, the exact fixed point is reached in just one step, irrespective of the size of the system. For more typical starting points, a few iterations are needed to reach a satisfactory accuracy. This method suppresses the scaling of the number of self-consistent cycles completely but has one extremely worrisome drawback: the computation of the preconditioning matrix scales asymptotically as order (N^4), with a large prefactor. Actually, the HIJ method is not at all used in state-of-the art computations, as the cost of computing the preconditioning matrix is prohibitive compared to the rest of the computation.

Sawamura and Kohyama[15] have proposed a method to decrease the prohibitive cost of HIJ. However, this algorithm still requires the use of the complete set of wave functions which is the main drawback of HIJ. Also, Auer and Krotscheck[17] have improved the proposal of HIJ by using the closure relation in order to limit the number of wave functions that must be explicitly included in the formulation of the dielectric matrix. However, the formulation they have proposed is limited to systems having no nonlocal potential.

Here, we propose a method inspired by HIJ. Similarly to (yet independently of) the method of Auer and Krotscheck,[17] it uses the closure relation, yet it remains applicable to all electronic structure calculation where wave functions are explicitly computed. In our method, several improvements over the original formulation enable tremendous improvement in applicability, thanks to a very low prefactor. The most time-consuming part of the HIJ algorithm is the computation of the susceptibility matrix, later to be combined with the Coulomb kernel and exchange-correlation kernel to generate the dielectric matrix. In order to generate the susceptibility matrix, HIJ relies on the Adler-Wiser expression[18,19] that contains a double sum over all occupied states and (a large number of) unoccupied states. We succeed to produce an *approximate* first-order preconditioner thanks to only occupied states (optionally completed by a few unoccupied states), with the closure relation providing an extrapolation of the susceptibility matrix to the exact one. Hence, we have nicknamed our procedure as the extrapolar method.

In HIJ the prefactor is also reduced by cutting off the short wave vector response of the system, thereby decreasing dramatically the size of the preconditioning matrix. However, throwing off such components completely destroys the preconditioning effect when applied to systems containing some vacuum. Indeed, the spurious small short-wavelength oscillations created by this approximation strongly influence the response of the exchange-correlation potential in the empty region of the cell. We show here that instead of the exact DFT preconditioner, the approximate RPA preconditioner (in which the exchange-correlation kernel is neglected) is free of this defect, still containing the potentiality to eliminate the scaling of the number of self-consistency cycles, and providing a large decrease of the number of iterations with respect to methods that do not rely on a preconditioner.

On this basis, extrapolar is applicable to any system, being metallic or insulating, homogeneous or inhomogeneous: it cuts down, for all systems, the number of self-consistent-field (SCF) iterations, often dramatically. Its overhead is reasonably small for systems with less than about 100 atoms, and extrapolar could even prove useful up to 1000 atoms in case an order (N^3) algorithm is used for solving the Kohn-Sham equation.

The first part (Sec. II) of this article gives the background needed to discuss preconditioning techniques. There, we explain the distinction we make between preconditioning and advanced mixing algorithms along with details on the conditioning problem, and the description of the HIJ preconditioning method. Sec. III describes our proposed preconditioner, whose parameter tuning is described in Sec. IV. Lastly, Sec. V is dedicated to evaluation of extrapolar generality and efficiency in various combinations with mixing algorithms. Atomic units (Hartree) are used throughout.

II. BACKGROUND

A. Self consistency

Within DFT, the Kohn-Sham equation determines the wave functions $\{\psi_i\}$ for a given, fixed, total external potential V_{ext}, and an input Hartree and exchange-correlation potential $V_{\text{Hxc,in}}$:

$$\left[-\frac{\nabla^2}{2} + V_{\text{ext}} + V_{\text{Hxc,in}}\right]|\psi_i\rangle = \varepsilon_i|\psi_i\rangle. \tag{1}$$

Such wave functions $\{\psi_i\}$ are orthonormalized $\langle\psi_i|\psi_j\rangle = \delta_{ij}$ and ε_i is the eigenenergy associated with the wave function ψ_i. The electronic density $\rho(\mathbf{r})$ is then determined from the wave functions,

$$\rho(\mathbf{r}) = \sum_i \psi_i^*(\mathbf{r})\psi_i(\mathbf{r}). \tag{2}$$

Finally, an output Hartree and exchange-correlation potential is determined from the density; the Hartree potential is a simple integral over all space of the density convoluted with the Coulomb kernel, and the exchange-correlation potential is defined as a functional derivative of the exchange-correlation energy with respect to the density:

$$V_{\text{Hxc,out}}(\mathbf{r}) = \int \frac{\rho(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|}d\mathbf{r}' + \frac{\delta E_{xc}}{\delta\rho(\mathbf{r})}. \tag{3}$$

Although the Hartree contribution is uniquely defined, there exist many different approximations for the exchange-correlation energy (whose exact form is not available). It will become apparent in the remaining of the paper that the

present approach does not rely on the precise definition of the exchange-correlation energy. However, some characteristics of this contribution will be revealed by the local-density approximation of the exchange-only part (LDAX), usually a dominating term in the exchange-correlation functional, for which

$$\frac{\delta E_{\mathrm{LDAX}}}{\delta \rho(\mathbf{r})} = \left(\frac{3}{\pi}\right)^{1/3} \rho(\mathbf{r})^{1/3}. \tag{4}$$

Altogether, the DFT [Eqs. (1)–(3)] treatment of the potential can be summarized by

$$V_{\mathrm{out}} = F[V_{\mathrm{in}}]. \tag{5}$$

(For sake of simplicity, we drop the Hxc label in the remaining of the paper.) The fixed point of this equation gives the self-consistent solution of the DFT equations

$$V^* = F[V^*]. \tag{6}$$

In order to find this solution the most widely used method is to guess a trial potential (alternatively one can start by guessing a trial density instead of a trial potential), then to determine the corresponding output potential, and from this knowledge, to make a better trial potential, also used to determine the corresponding output potential, and iterate this procedure. This process can be summarized by

$$V_{n,\mathrm{out}} = F[V_{n,\mathrm{in}}], \tag{7}$$

for n starting from 0, increasing with the number of available pairs of input and corresponding output potentials.

B. Linearization

First, let us suppose that we simply take the output potential as input potential

$$V_{n+1,\mathrm{in}} = V_{n,\mathrm{out}}. \tag{8}$$

A solution is reached when $V_{n,\mathrm{out}} \simeq V_{n,\mathrm{in}}$ within some prescribed tolerance. The main mathematical aspects which may prevent convergence are extensively discussed by Dederichs and Zeller in Ref. 6. If we write $\delta V = V - V^*$, Eq. (7) can be linearized

$$\delta V_{n,\mathrm{out}} = f \delta V_{n,\mathrm{in}}, \tag{9}$$

where f is the functional derivative of F. It is linked to the DFT electronic dielectric function, $\varepsilon_{\mathrm{DFT}}$, obtained from the independent-electron susceptibility of the system χ_0 and the Hartree and exchange-correlation Kernel K_{Hxc}:[11]

$$f = K_{Hxc}\chi_0 = 1 - \varepsilon_{\mathrm{DFT}}, \tag{10}$$

where 1 represents the unit matrix. At the fixed point of the self-consistency procedure that minimizes the total energy, one can prove that the DFT electronic dielectric function is definite positive.[29]

Now, we analyze the eigenvalues and eigenvectors of f; suppose P is a matrix such that $Pf_dP^{-1}=f$ and f_d is a diagonal matrix whose eigenvalues are λ_i, we get

$$\delta V_{n,\mathrm{out}} = P f_d^{n+1} P^{-1} \delta V_{0,\mathrm{in}}. \tag{11}$$

The convergence condition is now clear, the absolute value of the eigenvalues of f must all be lower than one: $|\lambda_i| < 1$. This constraint can be rephrased in term of the eigenvalues of the electronic dielectric function, $\varepsilon_{\mathrm{DFT},i} = 1 - \lambda_i$, that should all be positive (which is fulfilled) and smaller than two (which is not fulfilled for most solids and molecules).

Instead of using $V_{n+1,\mathrm{in}} = V_{n,\mathrm{out}}$, one can incorporate in $V_{n+1,\mathrm{in}}$ only part of the difference between $V_{n,\mathrm{out}}$ and $V_{n,\mathrm{in}}$, thanks to the simple mixing factor α (here a scalar):

$$\tilde{V}_{n+1,\mathrm{in}} = V_{n,\mathrm{in}} + \alpha(V_{n,\mathrm{out}} - V_{n,\mathrm{in}}). \tag{12}$$

In this case,

$$\delta V_{n,\mathrm{out}} = P[1 + \alpha(f_d - 1)]^{n+1} P^{-1} \delta V_{0,\mathrm{in}} \tag{13}$$

and convergence is reached if $|1+\alpha(\lambda_i - 1)| < 1$ or $|1 - \alpha\varepsilon_{\mathrm{DFT},i}| < 1$ for all eigenvalues. All the eigenvalues of $\varepsilon_{\mathrm{DFT}}$ being positive definite, it is enough to take $\alpha < 2/\max\{\varepsilon_{\mathrm{DFT},i}\}$ to make the simple mixing scheme converge. A more refined analysis[3] shows that the optimal rate of convergence of the simple mixing scheme (with the optimal mixing factor) depends on the ratio between the maximal and minimal eigenvalues, called the "condition number." The larger the condition number, the slower the convergence. If the condition number were equal to one, the convergence would be reached in just one step with the appropriate mixing factor. For metallic systems, the condition number diverges proportionally to the square of the largest linear dimension of the system. The reason for this divergence will become clear in Sec. III B.

C. Iterative schemes: Advanced mixing and preconditioning

Beyond simple mixing, two kinds of algorithms exist to help convergence. The "advanced mixing schemes" are widely spread (e.g., Anderson algorithm,[5] Pulay RMM-DIIS,[7] Dederichs[6] mixing and conjugate gradient[11]). Such algorithms try to guess the $V_{n+1,\mathrm{in}}$ from the ensemble of $(V_{\mathrm{in}}, V_{\mathrm{out}})$ pairs (or a subset of it) and eventually some other known properties, in the wisest possible way. On a quadratic problem with N eigenvalues they ensure convergence in, at most, N steps. For a problem where N is very large compared to the final number of iterations, for the best advanced mixing algorithms, the number of iterations will approximately scale as the square root of the condition number,[3] instead of the condition number. Moreover, since those algorithms take care of history, they can deal with evolutions of the Hessians with respect to convergence, that is, nonquadraticity. Their advantages over simple mixing actually apply irrespective of the f matrix definition.

Preconditioner algorithms are complementary to advanced mixing algorithms and can be used simultaneously. They are designed to directly modify eigenvalues and hence the condition number of the f matrix. Instead of using $V_{n,\mathrm{out}}$, one uses $\tilde{V}_{n,\mathrm{out}}$ which is the preconditioned potential, defined as follows:

$$\tilde{V}_{n,\mathrm{out}} = V_{n,\mathrm{in}} + \alpha(V_{n,\mathrm{out}} - V_{n,\mathrm{in}}), \tag{14}$$

with α being now a matrix, instead of a scalar as in the simple mixing algorithm. Then, we can mimic the analysis of the simple mixing algorithm (supposing no advanced mixing algorithm is used, for simplicity), and obtain

$$\delta\tilde{V}_{n+1,\mathrm{in}} = \delta\tilde{V}_{n,\mathrm{out}} = \delta V_{n,\mathrm{in}} + \alpha(\delta V_{n,\mathrm{out}} - \delta V_{n,\mathrm{in}}) \tag{15}$$

$$= [1 + \alpha(f-1)]\delta V_{n,\mathrm{in}} \tag{16}$$

$$= [1 + \alpha(f-1)]^{n+1}\delta V_{0,\mathrm{in}} \tag{17}$$

$$= [1 - \alpha\varepsilon_{\mathrm{DFT}}]^{n+1}\delta V_{0,\mathrm{in}}. \tag{18}$$

It is again possible to get a simple convergence criterion by an analysis of the eigenvalues of the operator $M_\alpha = 1 - \alpha\varepsilon_{\mathrm{DFT}}$ whose eigenvalues are μ_i. Convergence is ensured if $\lim_{n\to\infty} \mu_i^n = 0$. In case α is equal to the inverse of the electronic dielectric function, the eigenvalues of M_α are zero, hence the convergence is reached in just one step.

In the field of electronic structure calculations, the proposed preconditioners are: the HIJ approximation of the dielectric matrix,[9] Kerker's metal preconditioner[8] and some variants of it.[20,21]

D. The HIJ approach

The HIJ approach consists in computing first the susceptibility matrix thanks to the Adler-Wiser formula,[18,19] then computing exactly the dielectric matrix, and inverting it. Although the formula is written in a plane-wave basis in the original paper, with the number of allowed plane waves determined by a kinetic-energy cutoff, we will focus on the slightly simpler real-space formulation.

We suppose, also, for simplicity, that we deal with fully occupied or fully unoccupied bands only, and focus on the non-spin-polarized case. Here, n_{valence} is the number of occupied bands, which might vary with the wave vector k sampling the Brillouin zone. The susceptibility matrix writes

$$\chi_0(\mathbf{r},\mathbf{r}') = \sum_k \sum_{n=1}^{n_{\mathrm{valence}}} \sum_{n'=n_{\mathrm{valence}}+1}^{\infty} 2\frac{\psi_{kn}^*(\mathbf{r})\psi_{kn'}(\mathbf{r})\psi_{kn'}^*(\mathbf{r}')\psi_{kn}(\mathbf{r}')}{E_{kn'} - E_{kn}}. \tag{19}$$

Within the DFT formalism, the electronic dielectric matrix is given by:

$$\varepsilon_{\mathrm{DFT}}(\mathbf{r},\mathbf{r}') = \delta(\mathbf{r}-\mathbf{r}') - \int [K_H(\mathbf{r},\mathbf{r}'') + K_{xc}(\mathbf{r},\mathbf{r}'')]\chi_0(\mathbf{r}'',\mathbf{r}')d\mathbf{r}'', \tag{20}$$

where K_H is the Hartree kernel, $K_H(\mathbf{r},\mathbf{r}') = \frac{1}{|\mathbf{r}-\mathbf{r}'|}$, and K_{xc} the exchange-correlation kernel defined as $K_{xc}(\mathbf{r},\mathbf{r}') = \frac{\delta^2 E_{xc}}{\delta\rho(\mathbf{r})\delta\rho(\mathbf{r}')}$.

The resulting inverted dielectric matrix is applied to the residual of the potential $V_{n,\mathrm{out}} - V_{n,\mathrm{in}}$. If performed using the full set of plane waves and the full set of bands, it is the exact first-order preconditioner. Ho *et al.* note that the computation can be done on a fairly restricted plane wave basis which reduces tremendously the computational cost. Already at that time, the sum over unoccupied states was hindering the use of this method, although a full Hamiltonian diagonalization (thus giving all the bands, unoccupied as well as occupied) was then current practice. Actually, modern electronic structure calculations do not need anymore a full Hamiltonian diagonalization.

III. EXTRAPOLAR

A. Computing an approximate susceptibility matrix

The extrapolar method relies on a way for not including many conduction bands in the susceptibility matrix evaluation. In the following, we denote by n_{band} the number of bands explicitly included in the calculation, usually only slightly larger than the number of valence bands n_{valence}.

As seen from Eq. (19), it is possible to compute exactly the contribution to the susceptibility matrix of occupied and unoccupied bands up to n_{band}. We do it and denote the result as A.

$$A(\mathbf{r},\mathbf{r}') = \sum_k \sum_{n=1}^{n_{\mathrm{valence}}} \sum_{n'=n_{\mathrm{valence}}+1}^{n_{\mathrm{band}}} 2\frac{\psi_{kn}^*(\mathbf{r})\psi_{kn'}(\mathbf{r})\psi_{kn'}^*(\mathbf{r}')\psi_{kn}(\mathbf{r}')}{E_{kn'} - E_{kn}}. \tag{21}$$

Then:

$$\chi_0(\mathbf{r},\mathbf{r}') = A(\mathbf{r},\mathbf{r}') + B(\mathbf{r},\mathbf{r}') \tag{22}$$

where

$$B(\mathbf{r},\mathbf{r}') = \sum_k \sum_{n=1}^{n_{\mathrm{valence}}} \sum_{n'=n_{\mathrm{band}}+1}^{\infty} 2\frac{\psi_{kn}^*(\mathbf{r})\psi_{kn'}(\mathbf{r})\psi_{kn'}^*(\mathbf{r}')\psi_{kn}(\mathbf{r}')}{E_{kn'} - E_{kn}}. \tag{23}$$

We approximate part B, by replacing B with $\tilde{B}$ where the energy that depends on the summation index n' is replaced by a fixed energy $\bar{E}$, and the whole expression is multiplied by some parameter Λ:

$$B \simeq \tilde{B} = \Lambda \sum_k \sum_{n=1}^{n_{\text{valence}}} \sum_{n'=n_{\text{band}+1}}^{\infty} 2 \frac{\psi^*_{kn}(\mathbf{r})\psi_{kn'}(\mathbf{r})\psi^*_{kn'}(\mathbf{r}')\psi^*_{kn}(\mathbf{r}')}{\bar{E} - E_{kn}}. \tag{24}$$

Thanks to the closure relation, we can avoid the infinite summation. Λ helps in tuning the effect produced by replacing $E_{kn'}$ with $\bar{E}$. The final result is:

$$\tilde{B}(\mathbf{r},\mathbf{r}') = \Lambda \sum_k \sum_{n=1}^{n_{\text{valence}}} 2 \frac{\psi^*_{kn}(\mathbf{r})\psi_{kn}(\mathbf{r}')}{\bar{E} - E_{kn}} \times \left[\delta(\mathbf{r}-\mathbf{r}') - \sum_{n'=1}^{n_{\text{band}}} \psi_{kn'}(\mathbf{r})\psi^*_{kn'}(\mathbf{r}') \right] \tag{25}$$

$$= 2 \sum_k \sum_{n=1}^{n_{\text{valence}}} \frac{\Lambda}{\bar{E} - E_{kn}} \left[\rho_n(r)\delta(\mathbf{r}-\mathbf{r}') - \psi^*_{kn}(\mathbf{r})\psi_{kn}(\mathbf{r}') \sum_{n'=1}^{n_{\text{band}}} \psi_{kn'}(\mathbf{r})\psi^*_{kn'}(\mathbf{r}') \right]. \tag{26}$$

This enables a drastic reduction in the number of bands needed, still keeping close to the real susceptibility matrix.

Although the summation gets an important reduction, the overall scaling is unchanged since both A and B includes terms scaling such as $O(N^4)$. This makes the prefactor all the more important for actual use of this preconditioner.

B. Reciprocal space formulation

The complete approximated dielectric matrix ε_{DFT}, or Hessian of a DFT problem, can be a quite large object for large calculations, growing like N^2. Moreover, computing its inverse product with the potential could be a very long operation. Various ways exist to overcome this issue, both for approximated dielectric matrix[9] and for general Hessian.[12] As can be seen in Eq. (30), not all components of the susceptibility matrix will matter for preconditioning. In fact only those associated with the smallest **G** vectors will be amplified by the Coulombic term, resulting in the so-called charge sloshing phenomenon. This arises from the term $\frac{1}{\Omega_c \mathbf{G}^2}$ in Eq. (30) that produces the very large eigenvalues of some large Kohn and Sham SCF problem.

This led us to apply the method proposed by Ho *et al.* in Ref. 9. We compute the inverse dielectric matrix in reciprocal space, using only the components associated with **G** vectors smaller than an arbitrary plane-wave cutoff energy—usually much smaller than the cutoff energy used to describe the wave functions. The approximated susceptibility matrix in real-space writes

$$\chi_0(\mathbf{G},\mathbf{G}') = \sum_k \sum_{n=1}^{n_{\text{valence}}} \sum_{n'=n_{\text{valence}}}^{n_{\text{band}}} 2 \frac{\langle \psi_{kn'}|e^{-i(\mathbf{G}-\mathbf{G}')\mathbf{r}}|\psi_{kn'}\rangle\langle\psi_{kn}|e^{-i(\mathbf{G}-\mathbf{G}')\mathbf{r}}|\psi_{kn}\rangle}{E_{kn'} - E_{kn}} + \sum_k \sum_{n=1}^{n_{\text{valence}}} \frac{\lambda}{\bar{E} - E_{kn}} [\rho(\mathbf{G}-\mathbf{G}') - \langle\psi_{kn}|e^{i(\mathbf{G}-\mathbf{G}')\mathbf{r}}|\psi_{kn}\rangle \times \sum_{n'=1}^{n_{\text{band}}} \langle\psi_{kn'}|e^{-i(\mathbf{G}-\mathbf{G}')\mathbf{r}}|\psi_{kn'}\rangle] \tag{27}$$

In Sec. IV C, we mention that a typical cutoff energy used for extrapolar is about 2 hartrees. Usual soft pseudopotentials require cutoff energies of about 15 hartrees. The number of grid points used to represent a function with a given cutoff energy E is $N \propto \Omega_c E^{3/2}$. This makes the memory footprint of the susceptibility matrix grows as the square of the volume of the cell ($\propto \Omega_c^2 E^3$). The storage requirement for all wave functions also grows as the square of the volume of the cell, $\Omega_c^2 E^{3/2}$. The ratio of cutoffs and the fact that a single matrix is used results in a typical memory usage of a few percent of the one needed to store wave functions.

C. RPA vs DFT preconditioner

The procedure should now continue with the generation of the electronic dielectric matrix, see Eq. (20). However, a close look at the LDAX expression reveals a potential problem in its evaluation, in case of very weak densities:

$$\varepsilon_{\text{LDAX}}(\mathbf{r},\mathbf{r}') = \delta(\mathbf{r}-\mathbf{r}') - \int \frac{1}{|\mathbf{r}-\mathbf{r}''|} d\mathbf{r}'' \chi_0(\mathbf{r}'',\mathbf{r}') d\mathbf{r}'' - \frac{1}{3}\left(\frac{3}{\pi}\right)^{1/3} \rho^{-2/3}(\mathbf{r})\chi_0(\mathbf{r},\mathbf{r}'). \tag{28}$$

Indeed, the exchange-correlation kernel diverges in this case, like $\rho^{-2/3}(\mathbf{r})$. Because the susceptibility matrix is approximated by the treatment of the infinite sum over unoccupied bands and the reduction of the number of plane waves, this divergence is damaging and destroys the good preconditioning properties of the approximate scheme, as will be seen in Sec. VID.

In the RPA (or test-charge) dielectric matrix, the exchange and correlation part of Eq. (20) is removed, which gives:

$$\varepsilon_{\mathrm{RPA}}(\mathbf{r},\mathbf{r}') = \delta(\mathbf{r}-\mathbf{r}') - \int \frac{1}{|\mathbf{r}-\mathbf{r}''|} d\mathbf{r}'' \chi_0(\mathbf{r}'',\mathbf{r}') d\mathbf{r}''. \tag{29}$$

Its formulation in reciprocal space is especially instructive

$$\varepsilon_{\mathrm{RPA}}(G,\mathbf{G}') = \delta_{\mathbf{G},\mathbf{G}'} - \frac{1}{\Omega_c \mathbf{G}^2} \chi_0(\mathbf{G},\mathbf{G}'). \tag{30}$$

In the latter formula, there is another divergence, obtained when $\mathbf{G}=0$. However, the effect of this divergence is quite different from the one of the exchange term. Indeed, the range of eigenvalues of the electronic dielectric matrix, which gives the condition number of the self-consistency iterative procedure, is mainly governed by the magnitude of the smallest wave vector $\mathbf{G}$ treated,[3] through the Hartree term $\frac{1}{\Omega_c \mathbf{G}^2}$. For small values of $\mathbf{G}$, $\varepsilon_{\mathrm{RPA}}(\mathbf{G},\mathbf{G}')$ diverges for metals, while it saturates for insulators, due to the different behavior of $\chi_0(\mathbf{G},\mathbf{G}')$ in these limits.[22] On the contrary, the divergence in real space of the exchange-correlation kernel for vanishing densities is completely damped by the exact susceptibility matrix for vanishing densities, so that the divergence of the exchange-correlation term does not cause a divergence of the condition number. In general, the smaller eigenvalues are lower than one, but do not tend to zero. Thus, the RPA dielectric matrix might be used as a preconditioner as well.

Formulation of Eq. (29) has two advantages and one drawback compared to that of Eqs. (28) or (20). First and obviously, this approximation is exchange-correlation independent which is more tractable in a code; the associated drawback is, if divergence in the SCF cycles comes from the exchange and correlation term, it will not be treated. The second, yet most important and less obvious benefit is that the formulation of $\varepsilon_{\mathrm{RPA}}(\mathbf{G},\mathbf{G}')$ unlike $\varepsilon_{\mathrm{LDAX}}(\mathbf{G},\mathbf{G}')$ does not diverges in vacuum, which happens to be a strong problem when χ_0 is approximated by using a small plane wave cutoff energy. The standard use of extrapolar will be based on the RPA formulation.

IV. EXTRAPOLAR: IMPLEMENTATION AND TUNING

A. Implementation

The generation of the inverse dielectric matrix ought not to be done at each self-consistency step, since the overall goal is not to reach the exact computation of an approximate preconditioner, but to cut down the number of self-consistency cycles with the smaller possible overhead. The initialization of the self-consistency is usually done, either by considering a collection of pseudoatoms and the associated density and potential, or by a tight-binding calculation using atomic pseudo-orbitals. In both cases, the susceptibility matrix is already rather accurate when evaluated from the corresponding wave functions and eigenenergies. For instance, most of the time, its maximal eigenvalue will change by about a factor of two during SCF cycles. Moreover, it takes a very rough convergence in total energy to get a dielectric matrix perfectly converged for preconditioning purposes. Since its computation is long ($O(N^4)$ operation), one seeks a strategy for equilibration between preconditioner evaluation and advancement through SCF cycles. In the remaining of the paper, we will keep with a suboptimal, yet efficient scheme: the preconditioner is evaluated only at a few predetermined steps during SCF cycles, usually the first, and between the fourth and sixth step. This scheme proves efficient for most problems. However, it is advisable to use a wiser strategy for very large metallic systems.

In addition to this strategic choice, the extrapolar method depends on four parameters:

(i) the plane wave cutoff energy for the computation of χ_0 and ε;

(ii) the number of conduction bands explicitly included;

(iii) Λ;

(iv) $\bar{E}$.

We will have to tune these parameters in such a way that they can be applied to many different systems without further adjustment, giving a blind speedup.

The extrapolar technique has been implemented in the software package ABINIT[23] with the following default parameters: energy cutoff of 2.2 hartrees for evaluation of ε and similar matrices; $n_{\mathrm{bands}}=1.5n_{\mathrm{valence}}$ (for metals); $\Lambda=0.5$. $\bar{E}$ is taken to be the largest explicitly computed eigenenergy, plus 0.1 hartree. The default advanced mixing algorithm is Pulay RMM-DIIS algorithm. The following sections will provide an empirical justification of these choices, as well as highlight the impact (or lack of impact) they have. They have been tested for many more systems than those presented here, as this implementation in ABINIT has already been available for some years already. Concerning the strategy for the recomputation of the dielectric matrix, the default behavior recommended to the ingenuous user in ABINIT consists in a single update using the fifth step wave functions. Actually, on most systems, one to nine SCF steps are enough to get an excellent preconditioner (remember that even a factor of two on the eigenvalues is fine for preconditioning purposes). This works very well for all small systems ($L<10$ nm). For large systems, and specific mixing scheme, this recommendation must be tuned. Also, it is always useful to avoid recomputing the preconditioner at a SCF step where the potential, and then the updated wave functions, are further away to self-consistency than at the step where the preconditioner was previously computed.

B. Prototypical cells

Preconditioning problem arises from a combination of materials susceptibility and system size. In Ref. 3, Annett shows that in a worst case, a linear cell of bulk metal with a single long dimension, this leads to increase of eigenvalues proportional to the square of the system size. This worst case is most interesting for the present study, where we seek conditioning problems at the lowest CPU cost per self-consistent iteration, in order to tune parameters and study test cases. Yet, the preconditioner proposed by Kerker[8] is very efficient for bulk materials (or cells with weakly varying dielectric tensor). In order to tackle a real challenge, we will consider *inhomogeneous*, long, cells that simulate a periodic packing of bulk metal and vacuum layers (slabs), where Kerker's

TABLE I. Number of self-consistent-field cycles required to converge the total energy to 10^{-9} hartree, as a function of the plane-wave cutoff energy used to evaluate the preconditioner.

Cutoff [hartree]	2^2	2^1	2^0	2^{-1}	2^{-2}	2^{-3}	2^{-4}
$n_{\mathrm{SCFcycles}}$ for Sr S_8^4	20	20	20	21	17	19	∞
$n_{\mathrm{SCFcycles}}$ for Si S_4^2	22	21	24	24	21	25	∞

preconditioner is of nearly no help. We will base our nomenclature for the different slabs on the letter S (for Slab) with two integer indices, referring to multiples of primitive cell dimensions. The subscript index will mark out the number of crystal unit cell necessary to fill up the supercell and the superscript will mark out the number of such cells which are not empty. For instance S_{16}^8 designate a supercell simulating slabs wide of 8 unit cells separated by the same width of vacuum.

Strontium, a metal with a large lattice parameter of about 6 Å, described with a very soft pseudopotential using only two electrons, in $S_n^{n/2}$ cells, is prototypical of difficult conditioning problems, still leading to nonprohibitive calculations. Also, one of the most studied elements in DFT is silicon. Therefore, it will also be used within this study, despite being an insulator, with a maximal eigenvalue of the bulk electronic dielectric tensor on the order of twelve. Even though non-neutralized surfaces of silicon, with dangling bonds, will add some complexity to the problem, this last species of cells do not represent a tough conditioning problem. For each atomic species, the surface index chosen is [100] in their most stable phase under ambient conditions: FCC for Sr, diamond for Si. In both cases, the lattice parameter and internal geometric degrees of freedom of the slab are unrelaxed, as this is quite irrelevant for the present study of preconditioners (we draw general conclusions not only thanks to the present study of two kinds of cells, but also thanks to the extensive use of this preconditioner mentioned in Sec. IV A).

The pseudopotentials used here are the following: for strontium, the LDA and the GGA [PBE (Ref. 24)] pseudopotentials generated with FHI (Ref. 25) using the Troullier-Martins[26] scheme and provided on ABINIT website, both with two valence electrons; for silicon we used the Troullier-Martins[26] pseudopotential which has four valence electrons. Because the unit cell for strontium contains two atoms, the number of valence bands in Sr cells is just twice the number of filled unit cells. For instance we have eight valence bands in Sr S_8^4. In silicon, we have eight valence bands per filled unit cell; that is, for instance 16 in Si S_4^2.

We chose a quite shallow cutoff energy for plane waves of, respectively, 5 and 7 hartrees for strontium and silicon. This ensure a minimum absolute convergence of energy of 10^{-2} hartree. The integration in reciprocal space is made thanks to a tight k-points mesh. The minimum distance between them being about 30 bohrs^{-1} (corresponding, in case of bulk primitive cell, to a $2\times2\times2$ grid). Our tests on small cells show no difference in term of self-consistent field convergence between calculations with strongly converged parameters and those using lousily converged k-point mesh and energy cutoff. The later level of convergence allows faster tests of extrapolar on large simulation cells.

C. Tuning of parameters

In order to exemplify the behavior of extrapolar with respect to a variation of the four parameters we consider a S_8^4 cell of strontium and a S_4^2 cell of silicon, respectively.

We take as a reference the simple mixing approach without preconditioning nor advanced mixing algorithm, with mixing parameter varying from 1 to a very small value, by halving, and observe that the optimal mixing factor are 2^{-8} and 2^{-6}, respectively. For these choices, the total-energy convergence of 10^{-9} hartree is reached in 1713 and 454 steps, respectively.

When using extrapolar, we evaluate the preconditioner at step one and six only, use a mixing parameter $\alpha=0.4$, and $n_{\mathrm{bands}}=1.5n_{\mathrm{valence}}$. A small mixing factor is necessary to enforce convergence due to the absence of an advanced mixing algorithm. In this case, the total-energy convergence of 10^{-9} hartree is reached in 20 and 21 steps, respectively, for the strontium S_8^4 and the silicon S_4^2 cells. The speed-up factors with respect to simple mixing are thus on the order of 100 and 20, respectively. Notice that in Sec. V we examine long cells of strontium, for which the speed-up factor is quite larger.

Moreover, notice that the parameters chosen for Sec. IV C do not aim at optimum convergence. They are chosen to enlighten parameters effects.

1. Plane-wave cutoff energy

Table I presents the behavior of the number of steps as a function of the plane-wave cutoff. It shows that a very small cutoff energy for the preconditioner matrix, much smaller than the default value of 2.2 hartree, is possible. A smaller cutoff might even be a better choice with a simple mixing scheme based on a mixing parameter smaller than one (like here where it is 0.4). However, in most of the following, if not otherwise mentioned we keep the very safe value of 2.2 hartrees.

Considering that only small to very small wave vectors functions contribute to the Coulombic divergence of the self-consistency this result looks quite natural. For very large systems and when using advanced mixing schemes, it could even prove useful to use cutoff energies as low as 2^{-4} hartree taking into account that the divergence observed in Table I means simply that the largest eigenvalue of the SCF cycle, after preconditioning is larger than 5; something that all advanced mixing schemes can easily cope with.

2. Number of conduction bands

As demonstrated by Table II, that gives the behavior of the number of steps as a function of the number of conduction bands, the more conduction bands used, the faster is the

TABLE II. Number of SCF cycles required to converge the total energy to 10^{-9} hartree, as a function of the number of conduction bands used to evaluate the preconditioner matrix. Here, γ controls the number of bands via the following equation: $n_{\text{bands}}=n_{\text{valence}}\times(1+\gamma)$.

γ	2^1	2^0	2^{-1}	2^{-2}	2^{-3}
Sr conduction bands	16	8	4	2	1
$n_{\text{SCFcycles}}$ for Sr S_8^4	19	20	20	21	∞
Si conduction bands	32	16	8	4	2
$n_{\text{SCFcycles}}$ for Si S_4^2	20	20	21	22	24

convergence. However, this effect is quite weak, and not tremendously important when normal convergence rate is the goal, except for the smallest conduction bands number for strontium. In this case, the number of bands was not high enough for all the included electrons. $\gamma=0.5$, chosen for the other sections of this paper, looks like a very conservative and safe choice. If not really useful in the case of silicon it is mandatory for the longest strontium S cells because temperature spreads electrons very high in the conduction bands.

3. Λ and $\bar{E}$

In order to simplify inter systems comparisons, hereafter, numerical values given to $\bar{E}$ refers to the gap between the last included band in the calculation and the actual value of $\bar{E}$ in Eq. (24)–(26).

Evaluation of optimal values for the extrapolation parameters Λ and $\bar{E}$ can hardly be done in the previous mixing conditions; almost any choice for those parameters leads to good enough preconditioning for fast convergence. To discriminate between values, we use stringent conditions which result in fast divergence when the preconditioner is not good enough. Figure 1 shows an almost common valley of optimal choice for the parameters in both cells. To underline preconditioner effects, extrapolar was evaluated at each step and the mixing factors were raised to 0.8 for strontium S_8^4 and 1 for silicon S_4^2. Obviously, this figure keeps a large dependence in the other calculation parameters. Especially, the number of conduction bands explicitly included in the calculation. Nevertheless, values such as $\Lambda=0.5$ and $\bar{E}=0.1$ hartree seems a fair enough choice. More tests show no needs to change them whatever the supercell shape or the atomic species.

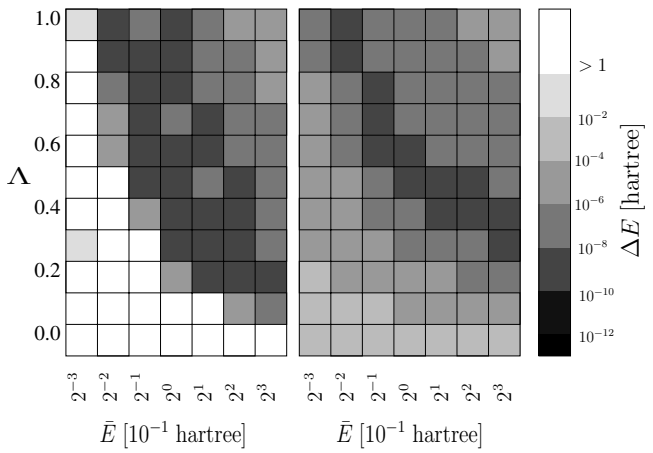

FIG. 1. Efficiency of extrapolar with respect to the choice of Λ and $\bar{E}$ in strontium (left) and silicon (right). The picture shows the convergence obtained after 10 SCF cycles in Sr (respectively, 6 in Si).

V. PERFORMANCES OF EXTRAPOLAR

In this section, we analyze the performances of the extrapolar-RPA algorithm, and show first its advantages with respect to the DFT version of extrapolar, then examine its behavior for larger cells. If not otherwise stated, parameters get their default value. For simple mixing we used $\alpha=0.5$.

A. RPA vs DFT preconditioner

In Fig. 2, the convergence using RPA and DFT formulation is considered for a bulk supercell and for a slab of strontium. For each case, the extrapolar operator is computed at startup and then at step 4 (denoted by the ticks labeled with *). The plane-wave cutoff energy is 1 hartree to underline the effect of reduced plane-wave cutoff when applying the DFT operator on cells including a vacuum part. A total of respectively 24 and 16 bands was used.

The DFT version leads to slightly better results in the bulk case (upper part of Fig. 2). This is coherent with the fact that

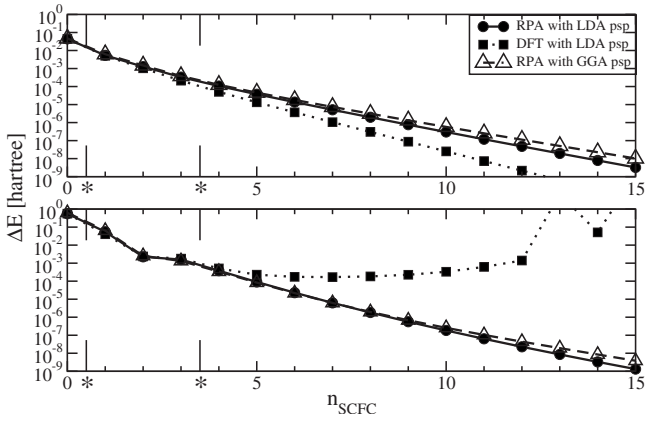

FIG. 2. Comparing the RPA and DFT versions of extrapolar on Sr S_8^8 and S_8^4. The dotted lines with squares show the evolution of total energy when using the DFT operator. The plain lines with circles and the dashed lines with triangles display the effect of the RPA version of extrapolar with, respectively, a LDA functional and a GGA functional.

TABLE III. Examples for the number of self-consistent-field cycles required to converge the total energy to 10^{-8} hartree as a function of iteration methodology and cell size: silicon and strontium.

	silicon				strontium			
Cell	S_4^2	S_8^4	S_{16}^8	S_{32}^{16}	S_8^4	S_{16}^8	S_{32}^{16}	S_{64}^{32}
Cell Length [Å]	21	43	86	172	48	97	194	389
n_{SCFC} (simple mixing+extrapolar)	15	14	14	15	14	15	14	18
n_{SCFC} (Pulay mixing only)	24	28	32	33	61	88	64	95
n_{SCFC} (Pulay mixing+extrapolar)	8	11	12	13	10	11	21	>60

the DFT preconditioner corresponds to the numerical approximation of the exact preconditioner, while the RPA preconditioner is the numerical approximation of an incomplete preconditioner. Yet its efficiency is plagued for the cell containing vacuum, as shown in the lower part of Fig. 2. The DFT preconditioner reevaluated between steps 3 and 4 does not work, although it ought to be more accurate than the RPA one, that perfectly does the job.

The DFT version of the extrapolar operator shall be avoided. In fact, the same thing hold for the HIJ preconditioner. As mentioned earlier, when used with a complete description no problem arises. Yet, when the operator is approximated either by reducing the plane-wave cutoff with respect to that of the pseudopotential or by means of extrapolar (reducing the number of bands explicitly included), calculations in cells including vacuum diverge. This is due to the exchange and correlation kernel which behaves proportionally to $\rho^{-2/3}$. When approximated, the divergence of this term prevents convergence.

Another advantage of the RPA version over its DFT counterpart is illustrated by the nearly joined convergence of the LDA and GGA systems when using the former. It works equally well irrespective of the exchange-correlation approximation, hereby reducing the implementation effort.

B. Condition number for large cells

The efficiency of a preconditioner can be quantified, in the neighborhood of the SCF solution, by the resulting condition number of the SCF cycles when this preconditioner is applied. For cells small enough, it is easy to evaluate it by a brute force approach. One makes first an almost perfectly converged evaluation of the wave functions, evaluate the preconditioning operator, and then impose a random perturbation to the potential. The norm of this perturbation must be small and kept constant through cycles. At convergence, the perturbation becomes the eigenvector corresponding to the maximally divergent eigenvalue of the system, while the normalizing factor used to keep a constant norm is this maximal eigenvalue. Using different mixing factor, one can obtain either the maximal or the minimal eigenvalue, whose ratio is the condition number.

Evaluation of the SCF condition number of nonpreconditioned Sr S_8^4, S_{12}^6, and S_{16}^8 cells, gives, respectively, 220, 500, and 790. This is very close to the expected behavior that ought to be quadratic with respect to the size of the system (the ratio between S_8^4 and S_{16}^8 is about 3.6 instead of 4). The preconditioned SCF condition numbers give 1.7, 1.6, and 1.4, respectively, two (nearly three) orders of magnitude smaller than the unpreconditioned condition numbers. So, if close to the solution, the extrapolar preconditioner achieves almost perfectly the goal of preconditioning: it makes SCF cycles convergence behave as $O(1)$ even for inhomogeneous, partly metallic systems.

C. Convergence for large cells

We exemplify now, for very long cells, up to S_{64}^{32}, the behavior of extrapolar and Pulay algorithms, each separately as well as in combination. Some typical results are displayed in Table III. Pulay algorithm belongs to the class of advanced mixing algorithms, and is often considered as the best one of this class.

Some important settings must be underlined. First, supercells used here are so long that finite numerical precision starts having noticeable effects on the total energy. This is why we reduced our convergence criterion on total energy to 10^{-8} hartree. Even this setting keeps having border effects for the largest cells. Second, in the ABINIT implementation of the Pulay algorithm, the default number of memorized input and output potential pairs is seven. This is not enough to reach convergence with the largest metallic cells, when Pulay mixing alone is used, so we have increased this number to fifteen, in this case only. Third, more evaluations of the dielectric matrix have been used, especially for strontium, which is much more difficult, from the conditioning point of view, than silicon. When used alone, extrapolar was evaluated every six steps (+ step 3 for strontium), after the initialization. Combined with the Pulay algorithm, extrapolar was evaluated at step 1 and 6 for silicon, respectively, 1, 3 and 6 for strontium. Fourth, because the evaluation time for the extrapolar operator has a fourth order dependence on system size, we chose to use a plane-wave cutoff energy smaller than usual, 0.5 Ha, to reduce the overhead of the preconditioner evaluation, and come close to the real situation. As shown previously, the impact of this choice is small.

As expected, Table III shows that energy convergence is no more dependent of system size when simple mixing is used. This can be opposed to the behavior of Pulay mixing alone which, as expected, requires an increasing number of steps to converge even when the condition number is not increased with cell size (this is the case of silicon where the condition number never exceeds twice the macroscopic permittivity). The faster convergence observed for silicon when

Pulay mixing is used in conjunction with extrapolar in the case of silicon was also expected.

Conversely, the effects of the combination of Pulay's mixing with extrapolar in the two largest strontium supercells may appear surprising. Especially when considering that, for instance, even evaluation every three steps of extrapolar do not prevent this very slow convergence. This bears witness to some interaction between advanced mixing scheme and the preconditioner. While the mixing scheme builds an implicit approximation of the Hessian of the system, each reevaluation of the preconditioner changes this Hessian leading to potentially harmful interaction. In the present case, the problem lies more likely in a bad schedule for the extrapolar evaluation; for large metallic supercells, Pulay mixing often leads to transient higher energy states. Evaluating extrapolar at one of these steps leads to a less accurate preconditioner. This is indeed happening for the largest simulations of strontium slabs. We are currently exploring scheduling strategies in which sufficient reduction in the potential residual is required before reevaluation of the preconditioner.

Due to these considerations, one must emphasize that Table III is not a systematic comparison between mixing methods and preconditioners (and their combined use). There are several purposes for this table. First, showing that the preconditioner indeed decreases the computational burden of the problem that initially had a wide spectrum of eigenvalue. Even with quite large cells, the convergence can be reached with the simplest mixing scheme. Second, to show typical increases of the number of SCF cycles with the size of the system: little increase of the number of SCF cycles in the case with preconditioner, and a regular increase of the number of SCF cycles in the Pulay case. Finally, to show one case where both techniques are joined, and comment about the difficulty found for Sr (as a warning).

D. Overhead

The overhead associated with the use of extrapolar can be decomposed in several contributions: a memory overhead discussed in Sec. III B, and a large computational one which decomposes as follows. First, the approximate susceptibility matrix must be evaluated. Then, the approximate RPA dielectric matrix is computed, and inverted to give the preconditioner matrix. Finally, the preconditioner matrix is applied to the potential residual. Only the last operation must be done at each self-consistent step. For large cells, the approximate susceptibility matrix evaluation is an $O(N^4)$ procedure, with a double sum over plane waves and a double sum over states; the computation of the RPA-dielectric matrix in reciprocal space, Eq. (30), is only an $O(N^2)$ procedure; the inversion of this matrix is an $O(N^3)$ operation; and the final application is an $O(N^2)$ procedure. The evaluation of the potential residual for an input potential is itself asymptotically an $O(N^3)$ procedure, to be done for each self-consistent step. Thus, for large cells, one expects a competition between the overhead due to the approximate matrix evaluation, only done a couple of times during the self-consistency, and the evaluation of the potential residual, to be performed at each self-consistent step. The inversion of the dielectric matrix might also count

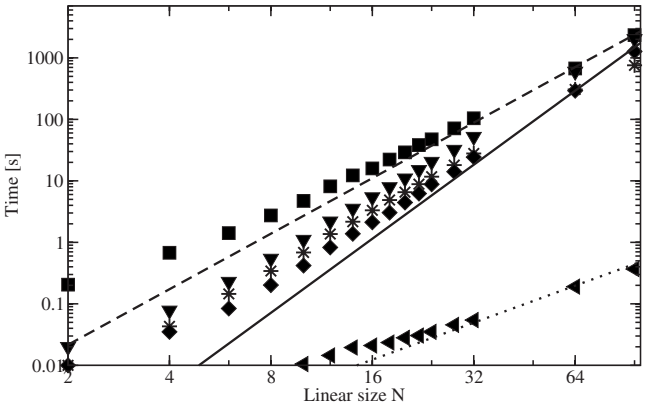

FIG. 3. Time scaling for simulations of Sr in cells $S_N^{N/2}$. Squares figure the time consumption for the wave-function optimization during a *single* SCF step, excluding all the operations needed for extrapolar. Triangles pointing down show the total time required to evaluate one extrapolar operator; this is split between the dielectric matrix inversion (stars) and the extrapolated susceptibility matrix evaluation (diamonds). Triangle pointing left stands for the time taken by the extrapolar application (time for one application). The continuous line is ax^4, the dotted line is bx^2, and the dashed line is cx^3.

in the global load of the calculation. Notice, however, that a more advanced implementation than the one described here could avoid the inversion of the dielectric matrix, by direct resolution of the linear system involving the potential residual and the dielectric matrix, thus replacing the $O(N^3)$ step to be done a few times by a $O(N^2)$ step to be done at each SCF cycle.

Compared to the HIJ algorithm, extrapolar simply reduces the prefactor of the global procedure. An accurate evaluation of the susceptibility matrix is done with a restricted set of plane waves, even for inhomogeneous cells. Reduction numbers are quite system dependent. They strongly depend on the ratio of plane-wave cutoff between the preconditioning operator and the overall system. The last improvement provided by extrapolar is related to the number of wave functions to be included in the preconditioner. For our small cells of strontium, S_8^4, this ratio between the number of plane waves and the number of bands is higher than 100. Then extrapolar provides a reduction of the calculation time to a single percent of HIJ requirements. The evaluation costs for the susceptibility matrix is proportional to the number of conduction bands used for a given system, and in our strontium studies, it decreases by a factor of about 10.

Figure 3 displays timings measurements related to extrapolar, as well as guidelines displaying power functions of order 2, 3, and 4. Extrapolar is evaluated with a plane-wave energy cutoff of 1 hartree. We have seen earlier that an even lower cutoff might be used, with very favorable decrease of the overhead due to extrapolar. For all system displayed in this picture, evaluating extrapolar takes always less time than doing *one* SCF iteration, while the cost of the application of the extrapolar preconditioner to the residual at each time step is always negligible. If we hypothesize that the number of self-consistent cycles to reach sufficient convergence is 15, and that extrapolar is evaluated twice, then, the overhead of

extrapolar for the S_{96}^{48} cell, with 48 atoms will only be about 10% of the total time. Under the same conditions, the overhead of extrapolar becomes comparable to the potential residual evaluation time for cells with about 300 atoms. For cells with this size, if one relies only on advanced mixing algorithms, like Pulay, the expected number of self-consistent cycles is at least one order of magnitude bigger than with extrapolar. Decreasing the cut off energy to a value lower than 1 hartree should raise the advantage of using extrapolar up to the 1000 atom level.

VI. CONCLUSION

In this article, we have proposed and analyzed extrapolar, an approximation of the HIJ exact preconditioner for self-consistently solving the DFT problem. This approximation allows to reduce drastically the computational cost associated with the use of the HIJ operator. In cases where the number of iterations increases with the system size (metallic cases), the extrapolar preconditioner allows to suppress this scaling, with sometimes order of magnitude speed up in the simulation. We showed that neglecting the exchange-correlation contribution to the dielectric matrix was an essential point for reaching wide applicability of this procedure in case of drastically reduced number of plane waves and conduction bands. The resulting extrapolar-RPA operator is sufficient to lead to independence of the number of cycles upon system size.

Our numerical experiments also show the efficiency of using both an advanced mixing algorithm and a preconditioner. However, this coupling appears to be more difficult for larger systems. This opens the question of development of nontrivial coupling algorithm between preconditioners and advanced mixing schemes. Finally, we must underline the main limitation of our extrapolar preconditioner. Even though it performs quite well, reducing dramatically the prefactor with respect to HIJ, its evaluation keeps scaling like $O(N^4)$. This reduces its applicability to system of less than a thousand of atoms.

ACKNOWLEDGMENTS

We thank S. Goedecker, T. Deutsch, D. Caliste, and M. Rayson for discussions about the present topics of research. We would like to acknowledge funding from the sixth European framework program, through the project BigDFT, contract NEST-2003-1 ADVENTURE 51181 5, the FRFC project 2.4502.05 "Simulation numérique. Application en physique de l'état solide océanographie et dynamique des fluides," the "Pôle d'attraction interuniversitaire," PAI/UIAP phase VI, "Quantum effects in clusters and nanowires," as well as the Communauté Française de Belgique, through the Action de Recherche Concertée "Nanosystemes hybrides metal-organiques."

[1] See, e.g., R. Martin, *Electronic Structure Basic Theory and Practical Methods* (Cambridge University Press, Cambridge, 2004), for instance.
[2] S. Goedecker, Rev. Mod. Phys. **71**, 1085 (1999).
[3] J. F. Annett, Comput. Mater. Sci. **4**, 23 (1995).
[4] D. Bowler, R. Choudhury, M. Gillan, and T. Miyazaki, Phys. Status Solidi B **243**, 989 (2006).
[5] D. G. Anderson, J. Assoc. Comput. Mach. **12**, 547 (1965).
[6] P. H. Dederichs and R. Zeller, Phys. Rev. B **28**, 5462 (1983).
[7] P. Pulay, Chem. Phys. Lett. **73**, 393 (1980).
[8] G. P. Kerker, Phys. Rev. B **23**, 3082 (1981).
[9] K.-M. Ho, J. Ihm, and J. D. Joannopoulos, Phys. Rev. B **25**, 4260 (1982).
[10] V. Eyert, J. Comput. Phys. **124**, 271 (1996).
[11] X. Gonze, Phys. Rev. B **54**, 4383 (1996).
[12] A. Sawamura, M. Kohyama, T. Keishi, and M. Kaji, Mater. Trans., JIM **40**, 1186 (1999).
[13] D. Bowler and M. Gillan, Chem. Phys. Lett. **325**, 473 (2000).
[14] A. Trellakis, A. T. Galick, A. Pacelli, and U. Ravaioli, J. Appl. Phys. **81**, 7880 (1997).
[15] A. Sawamura and M. Kohyama, Mater. Trans. **45**, 1422 (2004).
[16] J. Auer and E. Krotscheck, Comput. Phys. Commun. **118**, 139 (1999).
[17] J. Auer and E. Krotscheck, Comput. Phys. Commun. **151**, 265 (2003).
[18] S. Adler, Phys. Rev. **126**, 413 (1962).
[19] N. Wiser, Phys. Rev. **129**, 62 (1963).
[20] G. Kresse and J. Furthmüller, Phys. Rev. B **54**, 11169 (1996).
[21] Implemented within ABINIT (Ref. 23) since the late nineties by X. Gonze (unpublished).
[22] P. Ghosez, X. Gonze, and R. W. Godby, Phys. Rev. B **56**, 12811 (1997).
[23] X. Gonze *et al.*, Comput. Mater. Sci. **25**, 478 (2002); X. Gonze *et al.*, Z. Kristallogr. **220**, 558 (2005).
[24] J. P. Perdew, K. Burke, and M. Ernzerhof, Phys. Rev. Lett. **77**, 3865 (1996).
[25] Troullier-Martins (Ref. 27) type pseudopotentials made with FHI98PP (Ref. 28).
[26] Pseudopotentiels generated by D. C. Allan and A. Khein for ABINIT based on the method of Troullier and Martins (Ref. 27).
[27] N. Troullier and J. L. Martins, Phys. Rev. B **43**, 1993 (1991).
[28] M. Fuchs and M. Scheffler, Comput. Phys. Commun. **119**, 67 (1999).
[29] This follows from Eq. (16) in Ref. 11. Using the notation of this article we have $\frac{\delta^2 E^{trial}[v_{ext};u_{\min}]}{\delta u(\mathbf{r})\delta u(\mathbf{r}')}$ definite positive because it is taken at $u_{\min}$; $\chi_0(\mathbf{r},\mathbf{r}'')$ definite negative and then $\varepsilon_e(\mathbf{r}'',\mathbf{r}')$ definite positive.