SIAM J. SCI. COMPUT.
Vol. 35, No. 5, pp. S277–S298

ELLIPTIC PRECONDITIONER FOR ACCELERATING THE SELF-CONSISTENT FIELD ITERATION IN KOHN–SHAM DENSITY FUNCTIONAL THEORY*

LIN LIN† AND CHAO YANG†

Abstract. We discuss techniques for accelerating the self-consistent field iteration for solving the Kohn–Sham equations. These techniques are all based on constructing approximations to the inverse of the Jacobian associated with a fixed point map satisfied by the total potential. They can be viewed as preconditioners for a fixed point iteration. We point out different requirements for constructing preconditioners for insulating and metallic systems, respectively, and discuss how to construct preconditioners to keep the convergence rate of the fixed point iteration independent of the size of the atomistic system. We propose a new preconditioner that can treat insulating and metallic systems in a unified way. The new preconditioner, which we call an elliptic preconditioner, is constructed by solving an elliptic partial differential equation. The elliptic preconditioner is shown to be more effective in accelerating the convergence of a fixed point iteration than the existing approaches for large inhomogeneous systems at low temperature.

Key words. Kohn–Sham density functional theory, self-consistent field iteration, fixed point iteration, elliptic preconditioner

AMS subject classifications. 65F08, 65J15, 65Z05

DOI. 10.1137/120880604

1. Introduction.

Electron structure calculations based on solving the Kohn–Sham density functional theory (KSDFT) [23, 27] play an important role in the analysis of electronic, structural, and optical properties of molecules, solids, and other nanostructures. The Kohn–Sham equations define a nonlinear eigenvalue problem

$$\begin{aligned} &H[\rho]\psi_i = \varepsilon_i \psi_i, \\ &\rho(\mathbf{r}) = \sum_i f_i |\psi_i(\mathbf{r})|^2, \quad \int \psi_i^*(\mathbf{r})\psi_j(\mathbf{r})\,\mathrm{d}\mathbf{r} = \delta_{ij}, \end{aligned} \tag{1.1}$$

where the ε_i are the Kohn–Sham eigenvalues (or quasi-particle energies) and the ψ_i are called the Kohn–Sham wave functions or orbitals. These eigenfunctions define the electron density $\rho(\mathbf{r})$, which in turn defines the Kohn–Sham Hamiltonian

$$H[\rho] = -\frac{1}{2}\Delta + \mathcal{V}[\rho] + V_{\mathrm{ion}}, \tag{1.2}$$

where Δ is the Laplacian operator, $\mathcal{V}(\rho)$ is a nonlinear function of ρ, and V_{ion} is a potential function that is independent of ρ. The parameters $\{f_i\}$ that appear in the definition of ρ, which are often referred to as the occupation number, are defined by

$$f_i = \frac{1}{1+\exp(\beta(\varepsilon_i - \mu))}, \tag{1.3}$$

*Received by the editors June 11, 2012; accepted for publication (in revised form) May 8, 2013; published electronically October 28, 2013. This work was partially supported by the Laboratory Directed Research and Development Program of Lawrence Berkeley National Laboratory under the U.S. Department of Energy contract DE-AC02-05CH11231, and partially supported by Scientific Discovery through Advanced Computing (SciDAC) program funded by U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research and Basic Energy Sciences.
http://www.siam.org/journals/sisc/35-5/88060.html

†Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720 (linlin@lbl.gov, cyang@lbl.gov).

where β is proportional to the inverse of the temperature T and μ is called the chemical potential chosen to ensure that the f_i's satisfy

$$\sum_i f_i = \int \rho(\mathbf{r})\,\mathrm{d}\mathbf{r} = N \tag{1.4}$$

for a system that contains N electrons. The right-hand side of (1.3) is known as the Fermi–Dirac function evaluated at ε_i. When β is sufficiently large, the Fermi–Dirac function behaves like a step function that drops from 1 to 0 at μ (which lies between ε_N and ε_{N+1}). Spin degeneracy is omitted here for simplicity.

In this paper, we assume the Kohn–Sham system (1.1) is defined within the domain $\Omega = [0, L]^3$ with periodic boundary conditions, and the number of electrons N is proportional to the volume of the domain.

Because the eigenvalue problem (1.1) is nonlinear, it is often solved iteratively by a class of algorithms called *self-consistent field (SCF) iterations*. We will show in the following that the SCF iteration can be viewed as a fixed point iteration applied to a nonlinear system of equations defined in terms of the potential $\mathcal{V}$ that appears in (1.2) or the charge density ρ. The function evaluation in each step of the SCF iteration is relatively expensive. Hence, it is desirable to reduce the total number of SCF iterations by accelerating its convergence. Furthermore, we would like the convergence rate to be independent of the size of the system. In the past few decades, a number of acceleration schemes have been proposed [3, 40, 25, 13, 22, 30, 41, 4, 35]. However, none of the existing methods provide a satisfactory solution to the convergence issues to be examined in this paper, especially the issue of size dependency.

The purpose of the paper is twofold. First, we summarize a number of ways to accelerate the SCF iteration. Many of the schemes we discuss already exist in both the physics and the applied mathematics literature [29, 28, 4, 5, 51, 15, 42, 48]. We analyze the convergence properties of these acceleration schemes. In our analysis, we assume a good starting guess to the charge density or the potential is available. Such a starting guess is generally not difficult to obtain in real applications. As a result, the convergence of the SCF iteration can be analyzed through the properties of the Jacobian operator associated with the nonlinear map defined in terms of the potential or the density. Acceleration schemes can be developed by constructing approximations to the Jacobian or its inverse. These acceleration schemes can also be viewed as preconditioning techniques for solving a system of nonlinear equations.

It turns out that the SCF iteration exhibits quite different convergence behavior for insulating and metallic systems [16, 39]. These two types of systems are distinguished by the gap between ε_N and ε_{N+1} as the number of electrons N or, equivalently, the system size increases to infinity. For insulating systems,

$$\lim_{N\to\infty} E_g > 0, \tag{1.5}$$

where $E_g = \varepsilon_{N+1} - \varepsilon_N$, whereas for metallic systems, $\lim_{N\to\infty} E_g = 0$. Different accelerating (or preconditioning) techniques are required for insulating and metallic systems.

The second purpose of this paper is to propose a new framework for constructing a preconditioner for accelerating the SCF iteration. The preconditioner constructed under this framework, which we call the elliptic preconditioner, provides a unified treatment of insulating and metallic systems. It is effective for complex materials that contain both an insulating and a metallic component. This type of system is

considered to be difficult [41] for a standard Kohn–Sham solver, especially when the β parameter in (1.3) is relatively large (or the temperature is low).

The paper is organized as follows. In section 2, we introduce the fixed point iteration for solving the Kohn–Sham problem, and the simple mixing method as the simplest acceleration method. More advanced preconditioning techniques are discussed in section 3. In section 4, we discuss the convergence behavior of the acceleration methods for increasing system sizes. Based on these discussions, a new preconditioner called the elliptic preconditioner is presented in section 5. The performance of the elliptic preconditioner is compared to existing techniques for one-dimensional model problems and a realistic three-dimensional problem in section 6. We conclude and discuss future work in section 7.

In this paper, the Kohn–Sham orbitals $\{\psi_i\}$ are assumed to be in $H^1(\Omega)$. In practical calculations, they are discretized in a finite-dimensional space such as the space spanned by a set of plane waves. As a result, each operator corresponds to a finite-dimensional matrix. Our discussion in this paper is not restricted to any specific type of discretization of the Kohn–Sham orbitals. To simplify our discussion, we will not distinguish operators defined on $H^1(\Omega)$ from the corresponding matrices obtained from discretization unless otherwise noted. This applies to both differential and integral operators. Neither will we distinguish integral operators from their kernels. For example, we may simply denote $f(\mathbf{r}) = A[g](\mathbf{r}) \equiv \int A(\mathbf{r}, \mathbf{r}')g(\mathbf{r}')\,\mathrm{d}\mathbf{r}'$ by $f = Ag$ and represent the operator A by $A(\mathbf{r}, \mathbf{r}')$.

2. Fixed point iteration and simple mixing. It follows from the spectral theory that the charge density ρ defined in (1.1) can be written as

$$\rho(\mathbf{r}) = \left[I + e^{\beta(H[\rho]-\mu I)}\right]^{-1}(\mathbf{r}, \mathbf{r}). \tag{2.1}$$

Here $[\cdot](\mathbf{r}, \mathbf{r})$ denotes the diagonal elements of a matrix, and I is an identity operator. That is, $\rho(\mathbf{r})$ is the diagonal part of the Fermi–Dirac function evaluated at the Kohn–Sham Hamiltonian. The right-hand side of (2.1) defines a fixed point map from ρ to itself.

A similar fixed point map is also defined implicitly in terms of the potential $V = \mathcal{V}[\rho](\mathbf{r})$, where

$$\mathcal{V}[\rho](\mathbf{r}) = \int \frac{\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}\,\mathrm{d}\mathbf{r}' + V_{\mathrm{xc}}[\rho](\mathbf{r}), \tag{2.2}$$

where the first term corresponds to the electron-electron repulsion, and $V_{\mathrm{xc}}[\rho]$ is a nonlinear functional of ρ and is known as the *exchange-correlation* potential that accounts for many-body effects of the electrons. The choice of V_{xc} is not unique. A number of expressions are available in the physics literature [11, 38, 7, 31, 37]. However, for the purpose of this paper, we do not need to be concerned with the explicit form of V_{xc}. It should be noted that V_{xc} is often much smaller than the electron-electron repulsion term in magnitude. If the Dirac exchange [14] is used, V_{xc} often contains a term proportional to $\rho^{1/3}$.

It follows from (2.1) and (2.2) that ρ is implicitly a function of the potential V, which we will denote by $\rho = F(V)$. The analysis we present below and the acceleration strategies we propose are applicable to both the density fixed point map (2.1) and the potential fixed point map

$$V = \mathcal{V}[F(V)]. \tag{2.3}$$

Without loss of generality, we will focus on the potential fixed point map (2.3) in the rest of the paper. We remark that evaluating $\mathcal{V}[\rho]$ in (2.2) requires solving a Poisson equation, but the computation of $F(V)$ requires either diagonalizing the Hamiltonian $H[\rho]$ or approximating the Fermi–Dirac function of $H[\rho]$ directly [18]. Therefore, computing $F(V)$ is much more costly than computing $\mathcal{V}[\rho]$.

The simplest method for seeking the solution of (2.3) is the fixed point iteration. In such an iteration, we start from some input potential V_1, and iterate the following equation

$$V_{k+1} = \mathcal{V}\left[F(V_k)\right], \tag{2.4}$$

until (hopefully) the difference between V_{k+1} and V_k is sufficiently small.

When V_k is sufficiently close to the fixed point solution V^*, we may analyze the convergence of the fixed point iteration (2.4) by linearizing the function $\mathcal{V}[F(\cdot)]$ defined in (2.3) at V^*.

If we define $\delta V_k = V_k - V^*$, subtracting V^* from both sides of (2.4) and approximating $\mathcal{V}[F(V_k)]$ by its first-order Taylor expansion at V^* yields

$$\delta V_{k+1} \approx \left.\frac{\partial \mathcal{V}}{\partial V}\right|_{V=V^*} \delta V_k, \tag{2.5}$$

where $(\partial\mathcal{V}/\partial V)|_{V=V^*}$ is the Jacobian of $\mathcal{V}[F(V)]$ with respect to V evaluated at V^*.

It follows from the chain rule that

$$\frac{\partial \mathcal{V}}{\partial V} = \frac{\partial \mathcal{V}}{\partial \rho}\frac{\partial F}{\partial V}.$$

Taking the functional derivative of $\mathcal{V}[F(V)]$ given in (2.2) with respect to $\rho(\mathbf{r})$ yields

$$\frac{\partial \mathcal{V}}{\partial \rho}(\mathbf{r},\mathbf{r}') = \frac{1}{|\mathbf{r}-\mathbf{r}'|} + \frac{\partial V_{\mathrm{xc}}}{\partial \rho}(\mathbf{r},\mathbf{r}'), \tag{2.6}$$

where the terms in (2.6) represent kernels of integral operators evaluated at $\mathbf{r}$ and $\mathbf{r}'$. The first term on the right-hand side of (2.6) is the Coulomb kernel. It will be denoted by $v_c(\mathbf{r},\mathbf{r}')$ below. The second term is the functional derivative of the exchange-correlation potential with respect to ρ. It is often denoted by $K_{\mathrm{xc}}(\mathbf{r},\mathbf{r}')$ and is a Hermitian matrix.

In the physics literature, the functional derivative of F with respect to V is often referred to as the independent particle polarizability matrix, and denoted by $\chi(\mathbf{r},\mathbf{r}')$. At zero temperature, $\chi(\mathbf{r},\mathbf{r}')$ is given by the Adler–Wiser formula [1, 50]

$$\chi(\mathbf{r},\mathbf{r}') = 2\sum_{n=1}^{N}\sum_{m=N+1}^{\infty} \frac{\psi_n(\mathbf{r})\psi_m^*(\mathbf{r})\psi_n^*(\mathbf{r}')\psi_m(\mathbf{r}')}{\varepsilon_n - \varepsilon_m}, \tag{2.7}$$

where (ε_i,ψ_i), $i = 1,2,\ldots$, are the eigenpairs defined in (1.1). Note that χ is a Hermitian matrix, and is negative semidefinite since $\varepsilon_n \le \varepsilon_m$.

Equation (2.5) can be iterated recursively to yield

$$\delta V_{k+1} \approx \left(\frac{\partial \mathcal{V}}{\partial \rho}\chi\right)^k \delta V_1. \tag{2.8}$$

Therefore, a necessary condition that guarantees the convergence of the fixed point iteration is

$$\sigma\left(\frac{\partial \mathcal{V}}{\partial \rho}\chi\right) < 1,$$

where $\sigma(A)$ is the spectral radius of the operator (or matrix) A.

Unfortunately this condition is generally not satisfied as we will show later. However, a simple modification of the fixed point iteration can be made to overcome potential convergence failure as long as $\sigma(\frac{\partial \mathcal{V}}{\partial \rho}\chi)$ is bounded.

The modification takes the form

$$V_{k+1} = V_k - \alpha\left(V_k - \mathcal{V}\left[F(V_k)\right]\right), \tag{2.9}$$

where α is a scalar parameter. The updating formula given above is often referred to as simple mixing. When V_j is sufficiently close to V^*, the error propagation of the simple mixing scheme is

$$\delta V_{k+1} \approx \delta V_k - \alpha\left(I - \frac{\partial \mathcal{V}}{\partial \rho}\chi\right)\delta V_k. \tag{2.10}$$

Notice that $I - (\partial \mathcal{V}/\partial \rho)\chi$ is simply the Jacobian of the residual function $V - \mathcal{V}[F(V)]$ with respect to V. We will denote this Jacobian by J. Its value at V^* will be denoted by J_*. In the physics literature, this Jacobian is often referred to as a *dielectric* operator [1, 50], and denoted by ε. Furthermore, when $\frac{\partial \mathcal{V}}{\partial \rho}$ is positive definite, ε only has real eigenvalues because it can be symmetrized through a similarity transformation

$$\widetilde{\varepsilon} = \left(\frac{\partial \mathcal{V}}{\partial \rho}\right)^{-1/2} \varepsilon \left(\frac{\partial \mathcal{V}}{\partial \rho}\right)^{1/2} = I - \left(\frac{\partial \mathcal{V}}{\partial \rho}\right)^{1/2} \chi \left(\frac{\partial \mathcal{V}}{\partial \rho}\right)^{1/2},$$

where the symmetrized dielectric operator $\widetilde{\varepsilon}$ is Hermitian and has real eigenvalues. We remark that the assumption that $\frac{\partial \mathcal{V}}{\partial \rho}$ is positive definite may not always hold, especially when the material contains low electron density regions in which the exchange-correlation kernel K_{xc} contains large negative entries. However, because in general the product of K_{xc} and χ is much smaller in magnitude than $v_c\chi$, it is reasonable to expect that the eigenvalues of ε are close to those of $I - v_c\chi$, which are real. This type of approximation is also used in section 5 where we discuss how to construct an effective preconditioner for accelerating the fixed point iteration.

It follows from (2.10) that simple mixing will lead to convergence if

$$\sigma\left(I - \alpha J_*\right) < 1. \tag{2.11}$$

If $\lambda(J_*)$ is an eigenvalue of J_*, then the condition given in (2.11) implies that

$$|1 - \alpha\lambda(J_*)| < 1. \tag{2.12}$$

Consequently, $\lambda(J_*)$ must satisfy

$$0 < \alpha < \frac{2}{\lambda(J_*)}. \tag{2.13}$$

Note that (2.13) is only meaningful when $\lambda(J_*) > 0$ holds. Therefore, $\lambda(J_*) > 0$ is often referred to as the *stability condition* of a material [6, 33, 34]. Furthermore, when

$\lambda(J_*)$ is bounded, it is always possible to find a parameter α to ensure the convergence of the modified fixed point iteration even though the convergence may be slow.

We should comment that the stability condition $\lambda(J_*) > 0$ holds in most cases because $K_{\rm xc}\chi$ is typically much smaller in magnitude than $v_c\chi$. Note that v_c is positive definite, and χ is negative semidefinite. When the stability condition fails, phase transition may occur, such as the transition from uniform electron gas to Wigner crystals in the presence of low electron density [49]. Such a case is beyond the scope of the current study. Nonetheless, $\lambda(J_*)$ can become very large in practice even when the stability condition holds, especially for metallic systems of large sizes, as we will show in section 4. A large $\lambda(J_*)$ requires α to be set to a small value to ensure convergence. Even though convergence can be achieved, it may be extremely slow.

3. Preconditioned fixed point iteration and quasi-Newton acceleration. The simple mixing scheme selects α as a scalar in (2.9). If we replace the scalar α by the inverse of the Jacobian matrix of the function $V - \mathcal{V}[F(V)]$ evaluated at V_k, we obtain a Newton's update of V. When V_k is in the region where the linear approximation given by (2.5) is sufficiently accurate, Newton's method converges quadratically to the solution of (2.3).

3.1. Jacobian-free Krylov–Newton. The difficulty with applying Newton's method directly is that the Jacobian matrix $J_k = I - \frac{\partial \mathcal{V}[F(V)]}{\partial V}|_{V=V_k}$ or its inverse cannot be easily evaluated. However, we may apply a Jacobian-free Krylov–Newton technique [26] to obtain Newton's update

$$\Delta_k = J_k^{-1} r_k, \quad \text{where} \quad r_k = V_k - \mathcal{V}\left[F(V_k)\right],$$

by solving the linear system

$$J_k \Delta_k = r_k \tag{3.1}$$

iteratively using, for example, the GMRES algorithm [43]. The matrix vector multiplication of the form $y \leftarrow J_k x$, which is required in each GMRES iteration, can be approximated by finite difference

$$y \approx x - \frac{\mathcal{V}[F(V_k + \epsilon x)] - \mathcal{V}[F(V_k)]}{\epsilon}$$

for an appropriately chosen scalar ϵ.

The finite difference calculation requires one additional function evaluation of $\mathcal{V}[F(V_k+\epsilon x)]$ per GMRES step. Therefore, even though Newton's method may exhibit quadratic convergence, each Newton iteration may be expensive if the number of GMRES steps required to solve the correction equation (3.1) is large. The convergence rate of the GMRES method for solving the linear system (3.1) is known to satisfy [32]

$$\|\Delta_k^n - \Delta_k\| \leq C \left(\frac{\sqrt{\kappa(J_k)} - 1}{\sqrt{\kappa(J_k)} + 1} \right)^n \|\Delta_k^0 - \Delta_k\|, \tag{3.2}$$

where $\kappa(J_k) = \frac{\lambda_{\max}(J_k)}{\lambda_{\min}(J_k)}$ is the condition number of J_k, and Δ_k^n is the approximation of Δ_k at the nth step of the GMRES iteration. As we will show in section 4, the condition number $\kappa(J_k)$ can grow rapidly with respect to the size of the system, especially for metallic systems. Therefore the number of iterations required by an iterative solver also grows with respect to the size of the system unless preconditioning strategies are employed.

3.2. Broyden's and Anderson's methods. An alternative to Newton's method for solving (2.3) is a quasi-Newton method that replaces J_k^{-1} with an approximate Jacobian inverse C_k that is easy to compute and apply. In such a method, the updating strategy becomes

$$V_{k+1} = V_k - C_k \left(V_k - \mathcal{V}[F(V_k)]\right). \tag{3.3}$$

The simple mixing scheme discussed in the previous section can be viewed as a quasi-Newton method in which C_k is set to αI (or equivalently as a nonlinear version of the Richardson's iteration). More sophisticated quasi-Newton updating schemes can be devised by using Broyden's techniques [24] to construct better approximations to J_k or J_k^{-1}. In Broyden's second method, C_k is obtained by performing a sequence of low-rank modifications to some initial approximation C_0 of the Jacobian inverse using a recursive formula [15, 35] derived from the following constrained optimization problem:

$$\begin{aligned} &\min_{C} && \frac{1}{2}||C - C_{k-1}||_F^2 \\ &\text{s.t.} && S_k = CY_k, \end{aligned} \tag{3.4}$$

where C_{k-1} is the approximation to the Jacobian constructed in the $(k-1)$th Broyden iteration. The matrices S_k and Y_k above are defined as

$$S_k = (s_k, s_{k-1}, \ldots, s_{k-\ell}), \quad Y_k = (y_k, y_{k-1}, \ldots, y_{k-\ell}), \tag{3.5}$$

where s_j and y_j are defined by $s_j = V_j - V_{j-1}$ and $y_j = r_j - r_{j-1}$, respectively.

It is easy to show that the solution to (3.4) is

$$C_k = C_{k-1} + (S_k - C_{k-1}Y_k)Y_k^\dagger, \tag{3.6}$$

where $Y_k^\dagger$ denotes the pseudoinverse of Y_k, i.e., $Y_k^\dagger = (Y_k^T Y_k)^{-1} Y_k^T$. We remark that in practice $Y_k^\dagger$ is not constructed explicitly since we only need to apply $Y_k^\dagger$ to a residual vector r_k. This operation can be carried out by solving a linear least squares problem with appropriate regularization (e.g., through a truncated singular value decomposition).

A variant of Broyden's method is Anderson's method [3] in which C_{k-1} is fixed to an initial approximation C_0 at each iteration. It follows from (3.3) that Anderson's method updates the potential as

$$V_{k+1} = V_k - C_0(I - Y_k Y_k^\dagger)r_k - S_k Y_k^\dagger r_k. \tag{3.7}$$

In particular, if C_0 is set to αI, we obtain Anderson's method

$$V_{k+1} = V_k - \alpha(I - Y_k Y_k^\dagger)r_k - S_k Y_k^\dagger r_k.$$

commonly used in KSDFT solvers.

3.3. Pulay's method. An alternative way to derive Broyden's method is through a technique called direct inversion of iterative subspace (DIIS). The technique was originally developed by Pulay for accelerating the Hartree–Fock calculation [40]. Hence it is often referred to as *Pulay mixing* in the condensed matter physics community. The motivation of Pulay's method is to minimize the difference between V

and $\mathcal{V}[F(V)]$ within a subspace $\mathcal{S}$ that contains previous approximations to V. In Pulay's original work [40], the optimal approximation to V from $\mathcal{S}$ is expressed as $V_{\text{opt}} = \sum_{j=k-\ell-1}^{k} \alpha_j V_j$, where V_j $(j = k-\ell-1, \ldots, k)$ are previous approximations to V, and the coefficients α_j are chosen to satisfy the constraint $\sum_{j=k-\ell-1}^{k} \alpha_j = 1$.

When the V_j's are all sufficiently close to the solution of (2.3), $\mathcal{V}[F(\alpha_j V_j)] \approx \alpha_j \mathcal{V}[F(V_j)]$ holds. Hence we may obtain α_j (and consequently V_{opt}) by solving the following quadratic program

$$\begin{aligned} &\min_{\{\alpha_j\}} && \| \textstyle\sum_{j=k-\ell-1}^{k} \alpha_j r_j \|_2^2 \\ &\text{s.t.} && \textstyle\sum_{j=k-\ell-1}^{k} \alpha_j = 1, \end{aligned} \tag{3.8}$$

where $r_j = V_j - \mathcal{V}[F(V_j)]$.

Note that (3.8) can be reformulated as an unconstrained minimization problem if V_{opt} is required to take the form $V_{\text{opt}} = V_k + \sum_{j=k-\ell}^{k} \beta_j (V_j - V_{j-1})$, where β_j can be any unconstrained real number. Again, if we assume $\mathcal{V}[F(V)]$ is approximately linear at V_j and let $b = (\beta_{k-\ell}, \ldots, \beta_k)^T$, minimizing $\|V_{\text{opt}} - \mathcal{V}[F(V_{\text{opt}})]\|$ with respect to $\{\beta_j\}$ yields $b = -Y_k^{\dagger} r_k$, where Y_k is the same as that defined in (3.5).

In [28, 29], Pulay's method for updating V is defined as

$$V_{k+1} = V_{\text{opt}} - C_0 (V_{\text{opt}} - \mathcal{V}[F(V_{\text{opt}})]), \tag{3.9}$$

where C_0 is an initial approximation to the inverse of the Jacobian (at the solution). Substituting $V_{\text{opt}} = V_k - S_k Y_k^{\dagger} r_k$ into (3.9) yields exactly Anderson's updating formula (3.7).

3.4. Preconditioned fixed point iteration. If V_* is the solution to (2.3), then subtracting it from both sides of the quasi-Newton updating formula

$$V_{k+1} = V_k - C_k \left(V_k - \mathcal{V}[F(V_k)]\right)$$

yields

$$\delta V_{k+1} \approx \delta V_k - C_k J_* \delta V_k = (I - C_k J_*) \delta V_k, \tag{3.10}$$

where J_* is the Jacobian of the function $V - \mathcal{V}[F(V)]$ at V_* and C_k is the approximation to J_*^{-1} constructed at the kth step. If C_k is a constant matrix C for all k, we can rewrite (3.10) as

$$\delta V_{k+1} = (I - C J_*)^k \delta V_1. \tag{3.11}$$

Ideally, we would like to choose C to be J_*^{-1} to minimize the error in the linear regime. However, this is generally not possible (since we do not know V_*). However, if C is sufficiently close to J_*^{-1}, we may view C as a preconditioner for a preconditioned fixed point iteration defined by (3.11).

A desirable property for C is that $|\sigma(I - C J_*)| < 1$ or

$$0 < \sigma(C J_*) < 2. \tag{3.12}$$

Because $J_* = I - (\partial \mathcal{V} / \partial \rho)\chi$, we may construct C by seeking approximations to $\partial \mathcal{V} / \partial \rho$ and χ first and inverting the approximate Jacobian in (3.10). This is the approach taken by Ho, Ihm, and Joannopoulos in [22], which is sometimes known as the HIJ

approach. The HIJ approach approximates the matrix χ by using the Alder–Wiser formula given in (2.7) which requires computing all eigenvalues and eigenvectors associated with the Kohn–Sham Hamiltonian defined at V_k. The resulting computational cost for constructing χ alone is $\mathcal{O}(N^4)$ due to the explicit construction of each pair of $\psi_n(\mathbf{r})$ and $\psi_m(\mathbf{r})$ for $n = 1, 2, \ldots, N$ and $m = N+1, N+2, \ldots$. Such a preconditioning strategy is not practical for large problems.

An alternative to the HIJ approach is to use the "extrapolar" method proposed in [4]. This method replaces $\psi_m(\mathbf{r})$ by plane waves for large m. As a result, the number of ψ_m's that needs to be computed is reduced. However, such a reduction does not lead to a reduction in the computational complexity of constructing χ, which still scales as $\mathcal{O}(N^4)$. Therefore, the preconditioning strategy will become increasingly more expensive as the system size increases.

A more efficient preconditioner that works well for simple metallic systems is the Kerker preconditioner [25]. The potential updating scheme associated with this preconditioner is often known as the Kerker mixing scheme. The construction of the Kerker preconditioner is based on the observation that the Coulomb operator v_c can be diagonalized by the Fourier basis (plane waves), and the eigenvalues of the Coulomb operator are $4\pi/q^2$, where $q = |\mathbf{q}|$ is the magnitude of a sampled wave vector associated with the Fourier basis function of the form $e^{i\mathbf{q}\cdot\mathbf{r}}$. Furthermore, for simple metals the polarizability operator χ can be approximately diagonalized by the Fourier basis. The eigenvalues of χ are bounded from below and above. Therefore, if we omit the contribution from K_{xc}, the eigenvalues of J_* are $1 + 4\pi\gamma/q^2 = (q^2 + 4\pi\gamma)/q^2$ for some constant $\gamma > 0$ which is related to the Thomas-Fermi screening length [53]. But the true value of γ is generally unknown.

By neglecting the effect of K_{xc} in $\partial V/\partial\rho$, the Kerker scheme sets C to

$$C = \alpha\mathcal{F}^{-1}D_K\mathcal{F}, \tag{3.13}$$

where $\mathcal{F}$ is the matrix representation of the discretized Fourier basis that diagonalizes both v_c and χ, and the diagonal matrix D_K contains $q^2/(q^2+4\pi\hat{\gamma})$ on its diagonal, for some appropriately chosen constant $\hat{\gamma}$, and α is a parameter chosen to ensure (3.12) is satisfied.

As a result, the eigenvalues of CJ associated with the Kerker preconditioner are approximately $\alpha(q^2+4\pi\gamma)/(q^2+4\pi\hat{\gamma})$. When q is small, the corresponding eigenvalue of CJ is approximately $\alpha\gamma/\hat{\gamma}$. When q is large, the corresponding eigenvalue of CJ is approximately α. By choosing an appropriate $\alpha \in (0, 1)$ we can ensure that all eigenvalues are within $(0, 2)$ even when $\hat{\gamma}$ is not completely in agreement with the true γ.

The behavior of χ for simple insulating systems is very different from that for simple metallic systems. For simple insulating systems, the eigenvalues of χ corresponding to small q modes behave like $-\xi q^2$ where $\xi > 0$ is a constant [16, 39]. As a result, the spectral radius of J is bounded by a constant when the contribution from the exchange-correlation is negligible. Therefore, we can choose $C = \alpha I$ with an appropriate α to ensure the condition (3.12) is satisfied. The optimal choice of α will be discussed in the next section.

We should also note that when C_k is allowed to change from one iteration to another through the use of quasi-Newton updates, the convergence of the preconditioned fixed point (or quasi-Newton) iteration can be Q-superlinear [36].

4. Convergence rate and size dependency. In the previous section, we identified the condition under which a preconditioned fixed point iteration applied to the

Kohn–Sham problem converges. In this section, we discuss the optimal rate of convergence and its dependency on the size of the physical system. Ideally, we would like to construct a preconditioner to ensure the rate of convergence is independent of the system size.

4.1. The convergence rate of the simple mixing scheme. When the preconditioner is chosen to be $C = \alpha I$ (i.e., simple mixing), the convergence of the preconditioned fixed point iteration is guaranteed if α satisfies the condition given in (2.13). As is the case for analyzing Richardson's iteration for linear equations, it is easy to show that the optimal choice of α, which is the solution to the problem

$$r = \min_{\alpha} \max_{\lambda(J_*)} |1 - \alpha\lambda(J_*)|,$$

must satisfy

$$|1 - \alpha\lambda_{\max}| = |1 - \alpha\lambda_{\min}|, \tag{4.1}$$

where $\lambda_{\max}$ and $\lambda_{\min}$ are the largest and smallest eigenvalues of J_*, respectively.

The solution to (4.1) is

$$\alpha = \frac{2}{\lambda_{\max} + \lambda_{\min}}. \tag{4.2}$$

Therefore, the optimal convergence rate of simple mixing is simply [13]

$$r = \frac{\lambda_{\max} - \lambda_{\min}}{\lambda_{\max} + \lambda_{\min}} = \frac{\kappa(J_*) - 1}{\kappa(J_*) + 1}, \tag{4.3}$$

where $\kappa(J_*) = \lambda_{\max}/\lambda_{\min}$ is the condition number of the Jacobian at the solution.

4.2. The convergence rate of the Anderson/Pulay scheme. The convergence rate of Broyden's method can be shown to be Q-superlinear when it is applied to a smooth function, and when the starting guess of the solution is sufficiently close to the true solution and the starting guess of the Jacobian is sufficiently close to the true Jacobian at the solution [36]. However, in the Anderson or Pulay scheme, we reset the previous approximation to the Jacobian to $C_0 = \alpha I$ at each iteration. Therefore, its convergence may not be superlinear in general.

One interesting observation made by a number of researchers [2, 13] is that $V_{k+1} - V^*$ approximately lies in the Krylov subspace $\{V_0 - V^*, J_*(V_0 - V^*), \ldots, J_*^k(V_0 - V^*)\}$ when V_0 is sufficiently close to V^*, and V_{k+1} is constructed to have a minimum $\|V_{k+1} - V^*\|$ in this subspace in the Anderson/Pulay scheme even though we do not know this subspace explicitly (since we do not know V^* or J_*.) Therefore, one can draw a connection between the Anderson/Pulay scheme and the GMRES [43] algorithm for solving a linear system of equations [15, 42, 48]. As a result, if the Anderson or Pulay scheme converges, and when the pseudoinverse of Y_k is computed in exact arithmetic, heuristic reasoning suggests that the convergence rate may be approximately bounded by

$$r = \frac{\sqrt{\kappa(J_*)} - 1}{\sqrt{\kappa(J_*)} + 1}.$$

Clearly, when $\kappa(J_*)$ is large, the Anderson/Pulay acceleration scheme is superior to the simple mixing scheme.

When a good initial approximation to the inverse of the Jacobian (e.g., the Kerker preconditioner) C_0 is available, it can be combined with the Anderson/Pulay acceleration scheme to make the fixed point iteration converge more rapidly.

4.3. The dependency of the convergence rate on system size. A natural question that arises when we apply a preconditioned fixed point iteration to a large atomistic system is whether the convergence rate depends on the size of the system.

For periodic systems, the size of the system is often characterized by the number of unit cells in the computational domain. To simplify our discussion, we assume the unit cell to be a simple cubic cell with a lattice constant L. For nonperiodic systems such as molecules, we can construct a fictitious (cubic) supercell that encloses the molecule and periodically extend the supercell so that properties of the system can be analyzed through Fourier analysis. In both cases, we assume the number of atoms in each supercell is proportional to L^3.

Because the convergence rates of both the simple mixing and Anderson's method depend on the condition number of J_*, we should examine the dependency of $\kappa(J_*)$ with respect to L. When a good initial guess to the Jacobian C_0 is available, we should examine the dependency of $\kappa(C_0 J_*)$ with respect to L. Recall that $J_* = I - (v_c + K_{\rm xc})\chi$, where v_c is positive definite, $K_{\rm xc}$ is symmetric but not necessarily positive definite, and χ is symmetric negative semidefinite. The eigenvalues of J_* satisfy $\lambda(J_*) > \bar{\lambda} > 0$ where $\bar{\lambda}$ is independent of the system size. The inequality $\lambda_{\min} > \bar{\lambda} > 0$ gives the stability condition of the system.

The dependency of $\lambda_{\max}$ on L is generally difficult to analyze. However, for simple model systems such as a jellium system (or uniform electron gas) in which $K_{\rm xc}(\mathbf{r}, \mathbf{r}') = K^*_{\rm xc}\delta(\mathbf{r}, \mathbf{r}')$ for some constant $K^*_{\rm xc}$, we may use Fourier analysis to show that the eigenvalues of J_* are simply

$$\lambda_{\mathbf{q}} = 1 + \left(\frac{4\pi}{q^2} + K^*_{\rm xc}\right)\gamma F_L(q), \tag{4.4}$$

where $q = |\mathbf{q}|$, γ is a constant, and $F_L(q)$ is known as the *Lindhard* response function [53]. The Lindhard function satisfies

$$\lim_{q\to 0} F_L(q) = 1, \quad \lim_{q\to\infty} F_L(q) = 0. \tag{4.5}$$

Hence by taking $q = \frac{2\pi}{L}$, $\lambda_{\max}(J_*)$ is determined by $1 + \gamma(L^2/\pi + K^*_{\rm xc})$. As a result, the convergence of a fixed point iteration preconditioned by simple mixing and/or modified by Anderson's method tends to become slower for a jellium system as the system size increases.

When the Kerker preconditioner is used, the eigenvalues of CJ_* are

$$\lambda_{\mathbf{q}} = \alpha\frac{q^2 + \gamma F_L(q)(4\pi + K^*_{\rm xc}q^2)}{q^2 + 4\pi\hat{\gamma}}. \tag{4.6}$$

They are approximately α when q is large, and are determined by $\alpha F_L(q)\gamma/\hat{\gamma}$ when q is small. Since the smallest q satisfies $q = 2\pi/L$, the convergence rate of the Kerker preconditioned fixed point iteration is independent of system size for a jellium system. The same conclusion can be reached for simple metals such as Na or Al which behave like free electrons [53]. Therefore, the Kerker preconditioner is an ideal preconditioner for simple metals.

However, the Kerker preconditioner is not an appropriate preconditioner for insulating systems. Although in general the Jacobian associated with the insulating

system cannot be diagonalized by the Fourier basis, it can be shown that $e^{i\mathbf{q}\cdot\mathbf{r}}$ is an approximate eigenfunction of χ with the corresponding eigenvalue $-\xi q^2$ [16, 39]. If we neglect the contribution from $K_{\rm xc}$, $e^{i\mathbf{q}\cdot\mathbf{r}}$ is also an approximate eigenfunction of J_* with the corresponding eigenvalue $1+(4\pi/q^2)q^2\xi = 1+4\pi\xi$ for small q's. If C is chosen to be the Kerker preconditioner, then the corresponding eigenvalue of CJ_* is

$$\lambda_{\mathbf{q}} = \frac{q^2}{q^2+4\pi\hat{\gamma}}(1+4\pi\xi).$$

As the system size L increases, the smallest q, which satisfies $q = 2\pi/L$, becomes smaller. Consequently, the corresponding eigenvalue of CJ_* approaches zero. The convergence rate, which is determined by $\sigma(1-CJ_*)$, deteriorates as the system size increases.

For insulating systems, a good preconditioner is simply αI, where α is chosen to be close to $1/(1+4\pi\xi)$ (in general, we do not know the value of ξ). When such a preconditioner is used the convergence of the fixed point iteration becomes independent of the system size.

5. Elliptic preconditioner. As we have seen above, simple insulating and metallic systems call for different types of preconditioners to accelerate the convergence of a fixed point iteration for solving the Kohn–Sham problem. A natural question one may ask is how we should construct a preconditioner for a complex material that may contain both insulating and metallic components or metal surfaces.

Before we answer this question, we should point out that the analysis of the spectral properties of J_* that we presented earlier relies heavily on the assumption that the eigenfunctions of J_* are approximately plane waves. Although this assumption is generally acceptable for simple materials, it may not hold for more complex systems. Therefore, to develop a more general technique for constructing a good preconditioner, it may be more advantageous to explore ways to approximate J_*^{-1} or the solution to the equation $J_*\tilde{r}_k = r_k$ directly for some residue r_k.

One of the difficulties with this approach is in getting a good approximation of the polarizability operator χ in $J_* = I-(v_c+K_{\rm xc})\chi$. The use of the Adler–Wiser formula given in (2.7) would require computing almost all eigenpairs of H. Even when some of the ψ_m's can be replaced by simpler functions such as plane waves [4], constructing this operator and working with it would take at least $\mathcal{O}(N^4)$ operations.

Therefore it is desirable to replace the Adler–Wiser representation of χ with something much simpler and cheaper to compute. However, making such a modification to χ only may introduce undesirable error near low electron density regions because $K_{\rm xc}$ contains terms proportional to $\rho^{-2/3}$ (which originate from the Dirac exchange term [14]). This problem can be avoided by using the observations made in the physics community that the product of $K_{\rm xc}$ and χ is relatively small compared to $v_c\chi$, even though this observation has not been rigorously proved. As a result, it is reasonable to approximate J_* by $\tilde{J}_* = I - v_c\chi$. For historical reasons, this approximation is known as the random phase approximation (RPA) in the physics literature [4].

Note that, under RPA, we may rewrite $\tilde{J}_*^{-1}$ as

$$\tilde{J}_*^{-1} = (v_c^{-1}-\chi)^{-1}v_c^{-1}.$$

Since $v_c^{-1} = -\Delta/(4\pi)$, applying $\tilde{J}_*^{-1}$ to a vector r_k simply amounts to solving the following equation

$$(-\Delta - 4\pi\chi)\tilde{r}_k = -\Delta r_k. \tag{5.1}$$

To construct a preconditioner C, we will replace χ with a simpler operator. In many cases, we can choose the approximation to be a local (diagonal) operator defined by a function $b(\mathbf{r})$, although other types of more sophisticated operators are possible. To compensate for the simplification of χ, we replace the Laplacian operator on the left of (5.1) by $-\nabla\cdot(a(\mathbf{r})\nabla)$ for some appropriately chosen function $a(\mathbf{r})$. This additional change yields the following elliptic partial differential equation (PDE)

$$(-\nabla\cdot(a(\mathbf{r})\nabla)+4\pi b(\mathbf{r}))\,\tilde{r}_k=-\Delta r_k. \tag{5.2}$$

Because our construction of the preconditioner involves solving an elliptic equation, we call such a preconditioner an *elliptic preconditioner*.

Although the new framework we use to construct a preconditioner for the fixed point iteration is based on heuristics and certain simplifications of the Jacobian, it is consistent with the existing preconditioners that are known to work well with simple metals or insulators.

For example, for metallic systems, setting $a(\mathbf{r})=1$ and $b(\mathbf{r})=-\hat{\gamma}$ for some constant $\hat{\gamma}>0$ yields

$$(-\Delta+4\pi\hat{\gamma})\tilde{r}_k=-\Delta r_k. \tag{5.3}$$

The solution of the above equation is exactly the same as what is produced by the Kerker preconditioner.

For an isotropic insulating system, setting $a(\mathbf{r})=1+4\pi\xi$ and $b(\mathbf{r})=0$ yields

$$-(1+4\pi\xi)\Delta\tilde{r}_k=-\Delta r_k.$$

The solution to the above equation is simply

$$\tilde{r}_k=\frac{1}{1+4\pi\xi}r_k. \tag{5.4}$$

Such a solution corresponds to simple mixing with α set to $1/(1+4\pi\xi)$.

For a complex material that consists of both insulating and metallic components, it is desirable to choose approximations of $a(\mathbf{r})$ and $b(\mathbf{r})$ that are spatially dependent. The asymptotic behavior of χ with respect to the sizes of both insulating and metallic systems suggests that $a(\mathbf{r})$ and $b(\mathbf{r})$ should be chosen to satisfy $a(\mathbf{r})\geq 1$ and $b(\mathbf{r})\geq 0$. In this case, the operator defined on the left-hand side of (5.2) is a strongly elliptic operator. Such an operator is symmetric positive semidefinite.

The implementation of the elliptic preconditioner only requires solving an elliptic equation. In general $a(\mathbf{r}), b(\mathbf{r})$ are spatially dependent, and solving the elliptic preconditioner requires more than just a Fourier transform and scaling operation as is the case for the Kerker preconditioner. However, it is generally much less time consuming than evaluating the Kohn–Sham map or constructing χ or J_k. In particular, fast algorithms such as multigrid [8], fast multipole method (FMM) [20], hierarchical matrix [21] solver, and hierarchical semiseparable matrix (HSS) [12] can be applied to solve (5.2) with $\mathcal{O}(N)$ arithmetic operations. Even if we cannot achieve $\mathcal{O}(N)$ complexity, (5.2) can often be solved efficiently by Krylov subspaces iterative methods as we will show in the next section.

Our numerical experience suggests that simple choices of $a(\mathbf{r})$ and $b(\mathbf{r})$ can produce satisfactory convergence results for complicated systems. For example, if we place a metallic system in vacuum to ascertain its surface properties [44], we can

choose $b(\mathbf{r})$ to be a nonzero constant in the metallic region, and almost 0 in the vacuum part. The resulting piecewise constant function can be smoothed by convolving it with a Gaussian kernel. Similarly, $a(\mathbf{r})$ can be chosen to be 1 in the metallic region, and a constant larger than 1 in the vacuum region.

Due to the simplification that we made about the χ term in the Jacobian and the omission of the $K_{\mathrm{xc}}\chi$ term altogether, the construction of an elliptic preconditioner alone may not be sufficient to reduce the number of fixed point iterations required to reach convergence. However, such a preconditioner can be easily combined with the Broyden type of quasi-Newton technique such as Anderson's method discussed in 3.2 to further improve the convergence of the self-consistent field iteration. This is the approach we take in the examples that we will show in the next section.

6. Numerical results. In this section, we demonstrate the performance of the elliptic preconditioner proposed in the previous section, and compare it with other acceleration schemes through two examples. The first example consists of a one-dimensional (1D) reduced Hartree–Fock model problem that can be tuned to exhibit both metallic and insulating features. The second example is a three-dimensional (3D) problem we construct and solve in KSSOLV [52], which is a MATLAB toolbox for solving Kohn–Sham equations for small molecules and solids implemented entirely in MATLAB m-files. KSSOLV uses plane wave expansion to discretize the Kohn–Sham equations. It also uses the Troullier–Martins pseudopotential [47] to approximate the ionic potential.

6.1. One-dimensional reduced Hartree–Fock model. The 1D reduced Hartree–Fock model was introduced by Solovej [46], and has been used for analyzing defects in solids in [9, 10]. The simplified 1D model neglects the contribution of the exchange-correlation term. Nonetheless, typical behaviors of an SCF iteration observed for 3D problems can be exemplified by this 1D model. In addition to neglecting the exchange-correlation potential, we also use a pseudopotential to represent the electron-ion interaction. This makes our 1D model slightly different from that presented in [46].

The Hamiltonian in our 1D reduced Hartree–Fock model is given by

$$H[\rho] = -\frac{1}{2}\frac{d^2}{dx^2} + \int K(x,y)(\rho(y) + m(y))\,\mathrm{d}y. \tag{6.1}$$

Here $m(x) = \sum_{i=1}^{M} m_i(x - R_i)$, with the position of the ith nuclei denoted by R_i. Each function $m_i(x)$ takes the form

$$m_i(x) = -\frac{Z_i}{\sqrt{2\pi\sigma_i^2}} e^{-\frac{x^2}{2\sigma_i^2}}, \tag{6.2}$$

where Z_i is an integer representing the charge of the ith nucleus. The parameter σ_i represents the width of the nuclei in the pseudopotential theory. Clearly as $\sigma_i \to 0$, $m_i(x) \to -Z_i\delta(x)$ which is the charge density for an ideal nucleus. In our numerical simulation, we set σ_i to a finite value. The corresponding $m_i(x)$ is called a *pseudocharge density* for the ith nucleus. We refer to the function $m(x)$ as the total pseudocharge density of the nuclei. The system satisfies charge neutrality condition, i.e.,

$$\int \rho(x) + m(x)\,\mathrm{d}x = 0. \tag{6.3}$$

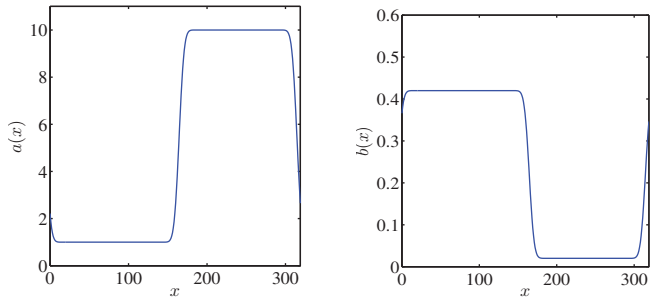

(a) (b)

FIG. 6.1. (a) *The choice of* $a(x)$ *and* (b) *the choice of* $b(x)$ *used by the elliptic preconditioner for a system with mixed metallic and insulating region.*

Since $\int m_i(x)\,\mathrm{d}x = -Z_i$, the charge neutrality condition (6.3) implies

$$\int \rho(x)\,\mathrm{d}x = \sum_{i=1}^{M} Z_i = N, \tag{6.4}$$

where N is the total number of electrons in the system. To simplify discussion, we omit the spin contribution here.

Instead of using a bare Coulomb interaction, which diverges in one dimension, we adopt a Yukawa kernel

$$K(x,y) = \frac{2\pi e^{-\kappa|x-y|}}{\kappa\epsilon_0}, \tag{6.5}$$

which satisfies the equation

$$-\frac{d^2}{dx^2}K(x,y) + \kappa^2 K(x,y) = \frac{4\pi}{\epsilon_0}\delta(x-y). \tag{6.6}$$

As $\kappa \to 0$, the Yukawa kernel approaches the bare Coulomb interaction given by the Poisson equation. The parameter ϵ_0 is used to make the magnitude of the electron static contribution comparable to that of the kinetic energy.

The parameters used in the reduced Hartree–Fock model are chosen as follows. Atomic units are used throughout the discussion unless otherwise mentioned. For all the systems tested below, the distance between each atom and its nearest neighbor is set to 10 a.u.. The Yukawa parameter $\kappa = 0.01$ is small enough that the range of the electrostatic interaction is sufficiently long, and ϵ_0 is set to 10.00. The nuclear charge Z_i is set to 2 for all atoms. Since spin is neglected, $Z_i = 2$ implies that each atom contributes to 2 occupied bands. The Hamiltonian operator is represented in a plane wave basis set. The temperature of the system is set to 100 K, which is usually considered to be very low, especially for the simulation of metallic systems.

By adjusting the parameters $\{\sigma_i\}$, the reduced Hartree–Fock model can be tuned to resemble an insulating, metallic, or hybrid system. We apply the elliptic preconditioner with different choices of $a(x)$ and $b(x)$ to all three cases. In the case of an insulator and a metal, both $a(x)$ and $b(x)$ are chosen to be constant functions. For the hybrid system, $a(x)$ and $b(x)$ are constructed by convolving a step function with a Gaussian kernel as shown in Figure 6.1. The σ_i values used for all these cases are listed in Table 6.1 along with the constant values chosen for $a(x)$ and $b(x)$ in the

TABLE 6.1
Test cases and SCF parameters used for the 1D model.

Case	σ_i	$a(x)$	$b(x)$	$\hat{\gamma}$
Insulating	2.0	1.0	0.0	0.50
Metallic	6.0	1.0	0.5	0.50
Hybrid	2.0/6.0	see Fig. 6.1(a)	see Fig. 6.1(b)	0.42

insulating and metallic cases. For the hybrid case, we partition the entire domain $[0, 320]$ into two subdomains, $[0, 160]$ and $[160, 320]$. The σ_i value is set to 6.0 in the first subdomain and 2.0 in the second subdomain.

For all three cases, we apply Anderson's method, Anderson's method combined with the Kerker preconditioner, and Anderson's method combined with the elliptic preconditioner to the SCF iteration. The α parameter used in the Anderson scheme is set to 0.50 in all tests. The $\hat{\gamma}$ parameter is set to 0.50 for the insulating and metallic cases, and 0.42 for the hybrid case.

The converged electron density ρ associated with the three 1D test cases as well as the 74 smallest eigenvalues associated with the Hamiltonian defined by the converged ρ are shown in Figure 6.2. The first 64 eigenvalues correspond to occupied states, and the rest correspond to the first 10 unoccupied states.

For the insulator case, the electron density fluctuates between 0.08 and 0.30. There is a finite gap between the highest occupied eigenvalue (ε_{64}) and the lowest unoccupied eigenvalue (ε_{65}). The band gap is $E_g = \varepsilon_{65} - \varepsilon_{64} = 0.067$ a.u.. The electron density associated with the metallic case is relatively uniform in the entire domain. The corresponding eigenvalues lie on a parabola (which is the correct distribution for uniform electron gas). In this case, there is no gap between the occupied eigenvalues and the unoccupied eigenvalues. For the hybrid case, the electron density is uniformly close to a constant in the metallic region (except at the boundary), and fluctuates in the insulating region. There is no gap between the occupied and unoccupied states.

In Figure 6.3, we show the convergence behavior of all three acceleration schemes for three test cases by plotting the relative self-consistency error in potential against the iteration number. In each one of the subfigures, the blue line with circles, the red line with triangles, and the black line with triangles correspond to tests performed on a 32-atom, 64-atom, and 128-atom system, respectively. We observe that the combination of Anderson's method and the elliptic preconditioner gives the best performance in all test cases. In particular, the number of SCF iterations required to reach convergence is more or less independent of the type of system and system size. We can clearly see that the use of the Kerker preconditioner leads to deterioration in convergence speed when the system size increases for insulating and hybrid systems. On the other hand, Anderson's method alone is not sufficient to guarantee the convergence of SCF iteration for metallic and hybrid systems. All these observed behaviors are consistent with the analysis we presented in the previous section.

6.2. Three-dimensional sodium system with vacuum. In this subsection, we compare the performance of different preconditioning techniques discussed in section 3 when they are applied to a 3D problem constructed in KSSOLV [52], a MATLAB toolbox for solving Kohn–Sham problems for molecules and solids. We have chosen to use the KSSOLV toolbox because of its ease of use, especially for prototyping new algorithms. The results presented here can be reproduced by other more advanced DFT software packages such as Quantum ESPRESSO [17], with some additional programming effort.

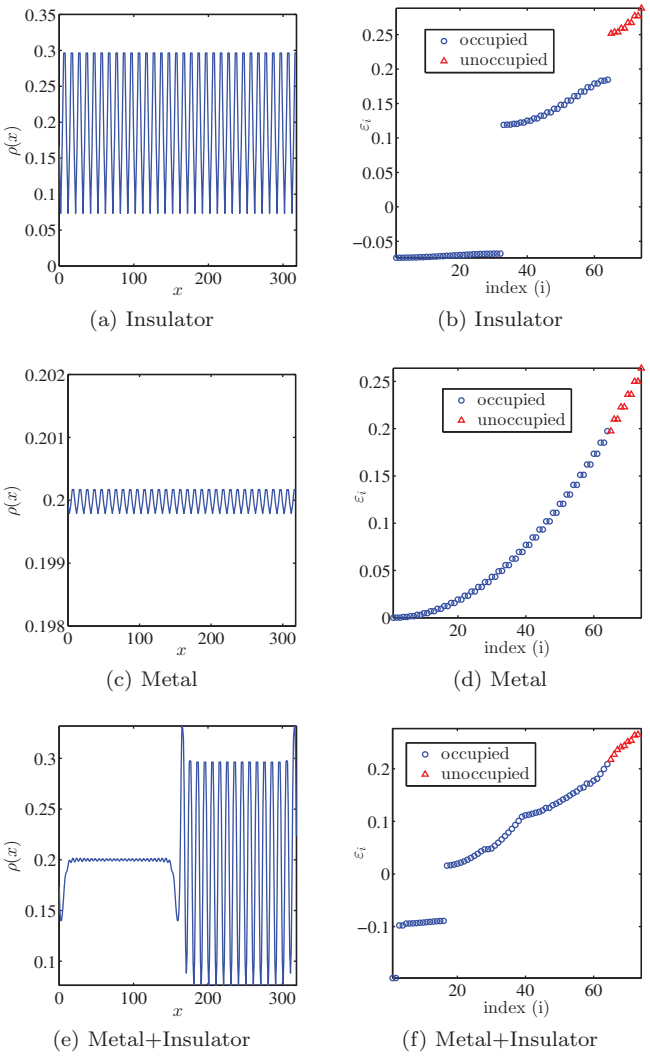

FIG. 6.2. *The electron density $\rho(x)$ of a* 32*-atom* (a) *insulating,* (c) *metallic, and* (e) *hybrid metal+insulator system in the left panel. The corresponding occupied (blue circles) and unoccupied eigenvalues (red triangles) are shown in the right panel in subfigures* (b), (d), (f), *respectively.*

The model problem we construct consists of a chain of sodium atoms placed in a vacuum region that extends on both ends of the chain. The sodium chain contains a number of body-centered cubic unit cells. The dimension of the unit cell along each direction is 8.0 a.u.. Each unit cell contains two sodium atoms. To examine the size dependency of the preconditioning techniques, we tested both a 16-unit cell (32-atom) model and a larger 32-unit cell (64-atom) model. The converged electron density on the $x = 0$ plane (or the [100] plane in crystallography terminology) associated with the 32-atom model is shown in Figure 6.4.

Figure 6.5 shows how Anderson's method, the combination of Anderson's method and the Kerker preconditioner, and the combination of Anderson's method and the elliptic preconditioner behave for both the 32-atom and the 64-atom sodium systems. For the 32-atom problem, the parameter α for the Anderson's method is set to 0.4.

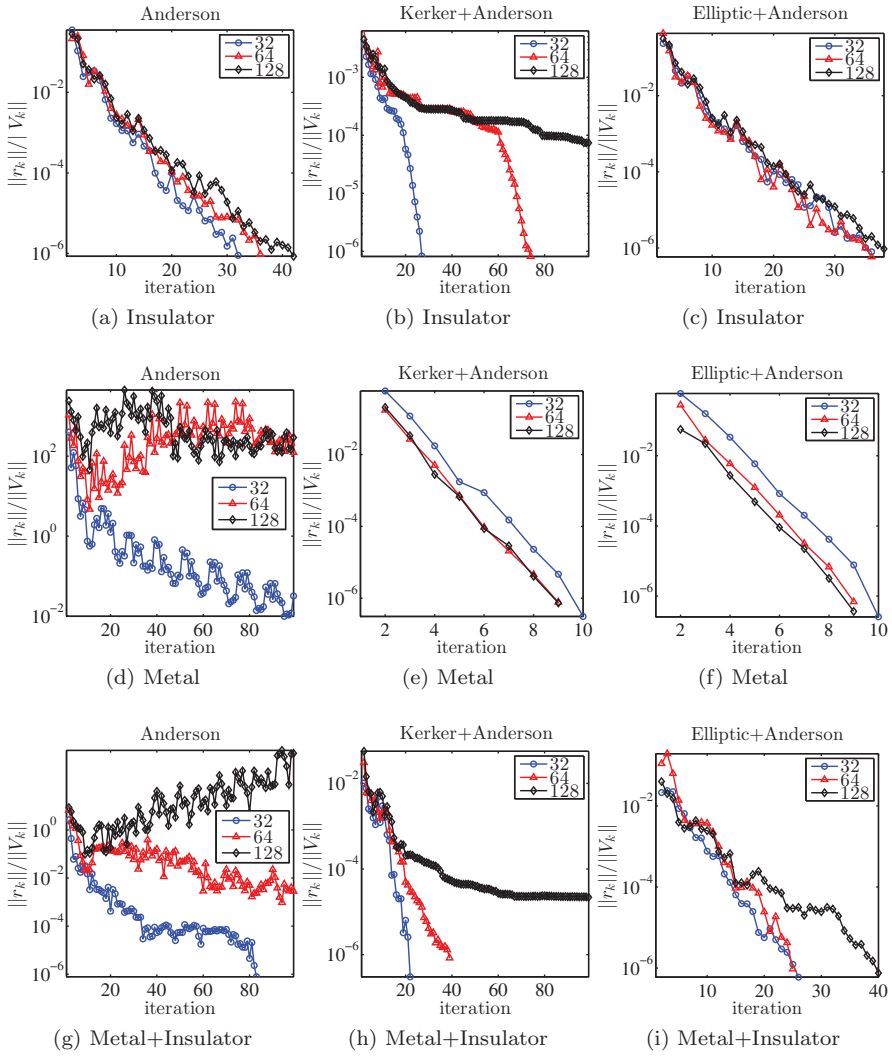

FIG. 6.3. *The convergence of Anderson's method, Anderson's method with the Kerker preconditioner, and Anderson's method with the elliptic preconditioner for insulators in subfigures* (a), (b), (c), *for metals in subfigures* (d), (e), (f), *and for hybrid systems in subfigures* (g), (h), (i), *respectively.*

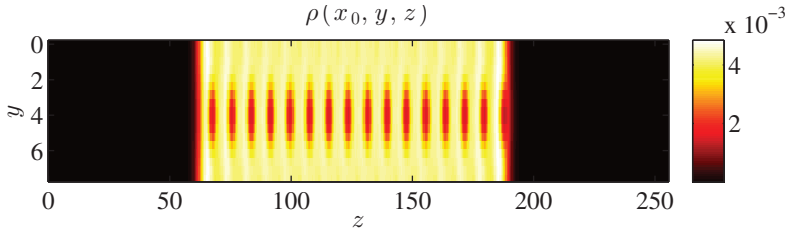

FIG. 6.4. *The $x = 0$ slice of the electron density $\rho(x, y, z)$ of the 32-atom sodium system with large vacuum regions at both ends.*

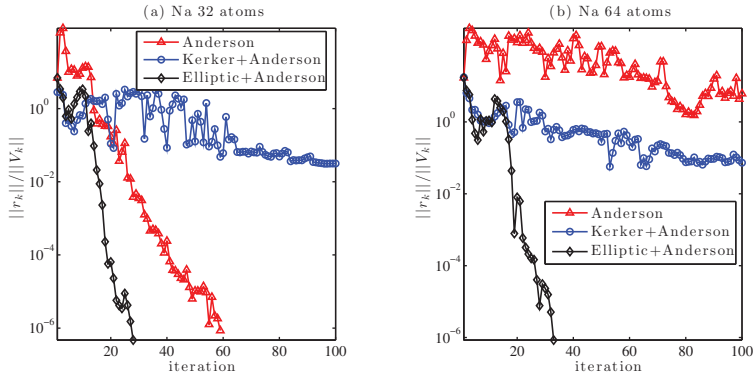

FIG. 6.5. *The convergence of Anderson's method, Anderson's method with the Kerker preconditioner, and Anderson's method with the elliptic preconditioner for quasi-1D Na systems with a large vacuum region with* 32 *Na atoms* (a) *and* 64 *Na atoms* (b).

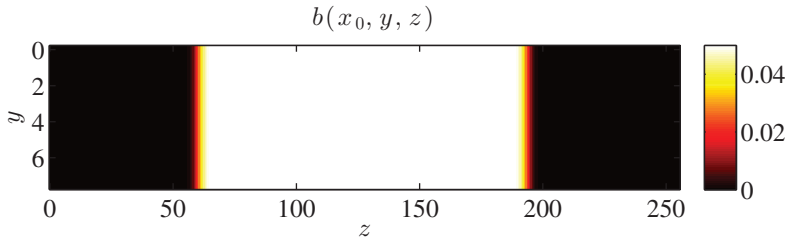

FIG. 6.6. *The* $x = 0$ *slices of the function* $b(x, y, z)$ *used in the elliptic preconditioner for a* 32*-atom sodium system with large vacuum regions at both ends.*

The parameter $\hat{\gamma}$ required in both the Kerker preconditioner and the elliptic preconditioner is set to 0.05. For the 64-atom problem, the parameter α is set to 0.8. For simplicity the function $a(x)$ required in the elliptic preconditioner is set to a constant function $a(x) = 1.0$. The $b(x)$ function (shown in Figure 6.6 for the 32-atom problem) is constructed by convolving a square wave function with a value of 0.05 in the sodium region and 0 in the vacuum region with a Gaussian kernel. The SCF iteration is declared to be converged when the relative self-consistency error in the potential is less than 10^{-6}.

As we can clearly see from Figure 6.5, the use of Anderson's method with the elliptic preconditioner leads to rapid convergence. Furthermore, the number of iterations (around 30) required to reach convergence does not change significantly as we move from the 32-atom problem to the 64-atom problem.

Using Anderson's method alone enables us to reach convergence in 60 iterations for the 32-atom problem. However, it fails to reach convergence within 100 iterations for the 64-atom case. When Anderson's method is combined with the Kerker preconditioner, the SCF iteration converges very slowly for both the 32-atom and the 64-atom problems.

7. Concluding remarks. We discussed techniques for accelerating the convergence of the self-consistent iteration for solving the Kohn–Sham problem. These techniques make use of the spectral properties of the Jacobian operator associated with the Kohn–Sham fixed point map. They can also be viewed as preconditioners

for a fixed point iteration. We pointed out the crucial difference between insulating and metallic systems and different strategies for constructing preconditioners for these two types of systems. A desirable property of the preconditioner is that the number of fixed point iterations is independent of the size of the system. We showed how this property can be maintained for both insulators and metals. Furthermore, we proposed a new preconditioner that treats insulating and metallic systems in a unified way. This preconditioner, which we refer to as an elliptic preconditioner, is constructed by solving an elliptic PDE with spatially dependent variable coefficients. Constructing preconditioners for insulating and metallic systems simply amounts to setting these coefficients to appropriate functions. The real advantage of this type of preconditioner is that it allows us to tackle more difficult problems that contain both insulating and metallic components at low temperature. We showed by simple numerical examples that this is indeed the case. Although the size of the systems used in our examples is relatively small because we are limited by the use of MATLAB, we can already see the benefit of an elliptic preconditioner in terms of keeping the number of SCF iterations relatively constant even as the system size gets larger. To fully test whether the preconditioner can achieve the goal of keeping the SCF iterations system size independent, we should implement the elliptic preconditioner in a standard electronic structure calculation software packages such as QUANTUM ESPRESSO [17], ABINIT [19], and SIESTA [45], etc., that are properly parallelized, which we plan to do in the near future.

Acknowledgments. We would like to thank Eric Cancès, Roberto Car, Weinan E, Weiguo Gao, Jianfeng Lu, Lin-Wang Wang, and Lexing Ying for helpful discussion.

REFERENCES

[1] S. L. Adler, *Quantum theory of the dielectric constant in real solids*, Phys. Rev., 126 (1962), pp. 413–420.
[2] H. Akai and P. H. Dederichs, *A simple improved iteration scheme for electronic structure calculations*, J. Phys. C, 18 (1985), pp. 2455–2460.
[3] D. G. Anderson, *Iterative procedures for nonlinear integral equations*, J. ACM, 12 (1965), pp. 547–560.
[4] P. M. Anglade and X. Gonze, *Preconditioning of self-consistent-field cycles in density-functional theory: The extrapolar method*, Phys. Rev. B, 78 (2008), 045126.
[5] J. F. Annett, *Efficiency of algorithms for Kohn-Sham density functional theory*, Comput. Mater. Sci., 4 (1995), pp. 23–42.
[6] R. Bauernschmitt and R. Ahlrichs, *Stability analysis for solutions of the closed shell Kohn–Sham equation*, J. Chem. Phys., 104 (1996), pp. 9047–9052.
[7] A. D. Becke, *Density-functional exchange-energy approximation with correct asymptotic behavior*, Phys. Rev. A (3), 38 (1988), pp. 3098–3100.
[8] A. Brandt, *Multi-level adaptive solutions to boundary-value problems*, Math. Comp., 31 (1977), pp. 333–390.
[9] E. Cancès, A. Deleurence, and M. Lewin, *A new approach to the modeling of local defects in crystals: The reduced Hartree-Fock case*, Comm. Math. Phys., 281 (2008), pp. 129–177.
[10] E. Cancès, A. Deleurence, and M. Lewin, *Non-perturbative embedding of local defects in crystalline materials*, J. Phys.: Condens. Matter, 20 (2008), 294213.
[11] D. M. Ceperley and B. J. Alder, *Ground state of the electron gas by a stochastic method*, Phys. Rev. Lett., 45 (1980), pp. 566–569.
[12] S. Chandrasekaran, M. Gu, and T. Pals, *A fast ULV decomposition solver for hierarchically semiseparable representations*, SIAM J. Matrix Anal. Appl., 28 (2006), pp. 603–622.
[13] P. H. Dederichs and R. Zeller, *Self-consistency iterations in electronic-structure calculations*, Phys. Rev. B, 28 (1983), pp. 5462–5472.
[14] P. A. M. Dirac, *On the theory of quantum mechanics*, R. Soc. London Proc., Ser. A Math. Phys. Eng. Sci., 112 (1926), pp. 661–677.
[15] H-R. Fang and Y. Saad, *Two classes of multisecant methods for nonlinear acceleration*, Numer. Linear Algebra Appl., 16 (2009), pp. 197–221.

[16] P. Ghosez, X. Gonze, and R. W. Godby, *Long-wavelength behavior of the exchange-correlation kernel in the Kohn-Sham theory of periodic systems*, Phys. Rev. B, 56 (1997), pp. 12811–12817.
[17] P. Giannozzi et al., *QUANTUM ESPRESSO: A modular and open-source software project for quantum simulations of materials*, J. Phys.: Condens. Matter, 21 (2009), 395502.
[18] S. Goedecker, *Linear scaling electronic structure methods*, Rev. Mod. Phys., 71 (1999), p. 1085.
[19] X. Gonze Cote, et al., *Abinit: First-principles approach to material and nanosystem properties*, Comput. Phys. Commun., 180 (2009), pp. 2582–2615.
[20] L. Greengard and V. Rokhlin, *A fast algorithm for particle simulations*, J. Comput. Phys., 73 (1987), pp. 325–348.
[21] W. Hackbusch, *A sparse matrix arithmetic based on $\mathcal{H}$-matrices. Part* I: *Introduction to $\mathcal{H}$-matrices*, Computing, 62 (1999), pp. 89–108.
[22] K. M. Ho, J. Ihm, and J. D. Joannopoulos, *Dielectric matrix scheme for fast convergence in self-consistent electronic-structure calculations*, Phys. Rev. B, 25 (1982), pp. 4260–4262.
[23] P. Hohenberg and W. Kohn, *Inhomogeneous electron gas*, Phys. Rev., 136 (1964), pp. B864–B871.
[24] D. D. Johnson, *Modified Broyden's method for accelerating convergence in self-consistent calculations*, Phys. Rev. B, 38 (1988), pp. 12807–12813.
[25] G. P. Kerker, *Efficient iteration scheme for self-consistent pseudopotential calculations*, Phys. Rev. B, 23 (1981), pp. 3082–3084.
[26] D. A. Knoll and D. E. Keyes, *Jacobian-free Newton–Krylov methods: A survey of approaches and applications*, J. Comput. Phys., 193 (2004), pp. 357–397.
[27] W. Kohn and L. Sham, *Self-consistent equations including exchange and correlation effects*, Phys. Rev., 140 (1965), pp. A1133–A1138.
[28] 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 (1996), pp. 15–50.
[29] G. Kresse and J. Furthmüller, *Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set*, Phys. Rev. B, 54 (1996), pp. 11169–11186.
[30] K. N. Kudin, G. E. Scuseria, and E. Cancès, *A black-box self-consistent field convergence algorithm: One step closer*, J. Chem. Phys., 116 (2002), pp. 8255–8261.
[31] C. Lee, W. Yang, and R. G. Parr, *Development of the Colle-Salvetti correlation-energy formula into a functional of the electron density*, Phys. Rev. B, 37 (1988), pp. 785–789.
[32] J. Liesen and P. Tichy, *Convergence analysis of Krylov subspace methods*, GAMM-Mitt. Ges. Angew. Math. Mech., 27 (2004), pp. 153–172.
[33] J. Lu and W. E, *Electronic structure of smoothly deformed crystals: Cauchy-born rule for the nonlinear tight-binding model*, Comm. Pure Appl. Math., 63 (2010), pp. 1432–1468.
[34] J. Lu and W. E, *The electronic structure of smoothly deformed crystals: Wannier functions and the Cauchy–Born rule*, Arch. Ration. Mech. Anal., 199 (2011), pp. 407–433.
[35] L. D. Marks and D. R. Luke, *Robust mixing for ab initio quantum mechanical calculations*, Phys. Rev. B, 78 (2008), 075114.
[36] J. Nocedal and S. J. Wright, *Numerical Optimization*, Springer-Verlag, New York, 1999.
[37] J. P. Perdew, K. Burke, and M. Ernzerhof, *Generalized gradient approximation made simple*, Phys. Rev. Lett., 77 (1996), pp. 3865–3868.
[38] J. P. Perdew and A. Zunger, *Self-interaction correction to density-functional approximations for many-electron systems*, Phys. Rev. B, 23 (1981), pp. 5048–5079.
[39] R. M. Pick, M. H. Cohen, and R. M. Martin, *Microscopic theory of force constants in the adiabatic approximation*, Phys. Rev. B, 1 (1970), pp. 910–920.
[40] P. Pulay, *Convergence acceleration of iterative sequences: The case of SCF iteration*, Chem. Phys. Lett., 73 (1980), pp. 393–398.
[41] D. Raczkowski, A. Canning, and L. W. Wang, *Thomas-Fermi charge mixing for obtaining self-consistency in density functional calculations*, Phys. Rev. B, 64 (2001), 121101.
[42] T. Rohwedder and R. Schneider, *An analysis for the DIIS acceleration method used in quantum chemistry calculations*, J. Math. Chem., 49 (2011), pp. 1889–1914.
[43] Y. Saad and M. H. Schultz, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869.
[44] D. S. Sholl and J. A. Steckel, *Density Functional Theory: A Practical Introduction*, Wiley-Interscience, Hoboken, 2009.
[45] J. M. Soler, E. Artacho, J. D. Gale, A. García, J. Junquera, P. Ordejón, and D. Sánchez-Portal, *The SIESTA method for ab initio order-N materials simulation*, J. Phys.: Condens. Matter, 14 (2002), pp. 2745–2779.
[46] J. P. Solovej, *Proof of the ionization conjecture in a reduced Hartree-Fock model*, Invent. Math., 104 (1991), pp. 291–311.

[47] N. Troullier and J. L. Martins, *Efficient pseudopotentials for plane-wave calculations*, Phys. Rev. B, 43 (1991), pp. 1993–2006.
[48] H. F. Walker and P. Ni, *Anderson acceleration for fixed-point iterations*, SIAM J. Numer. Anal., 49 (2011), pp. 1715–1735.
[49] E. Wigner, *On the interaction of electrons in metals*, Phys. Rev., 46 (1934), pp. 1002–1011.
[50] N. Wiser, *Dielectric constant with local field effects included*, Phys. Rev., 129 (1963), pp. 62–69.
[51] C. Yang, W. Gao, and J. C. Meza, *On the convergence of the self-consistent field iteration for a class of nonlinear eigenvalue problems*, SIAM J. Matrix Anal. Appl., 30 (2009), pp. 1773–1788.
[52] C. Yang, J. C. Meza, B. Lee, and L. W. Wang, *KSSOLV–a MATLAB toolbox for solving the Kohn–Sham equations*, ACM Trans. Math. Software, 36 (2009), p. 10.
[53] J. M. Ziman, *Principles of the Theory of Solids*, Cambridge University Press, New York, 1979.