Computer Physics Communications 184 (2013) 2478–2488

Contents lists available at ScienceDirect

Computer Physics Communications

journal homepage: www.elsevier.com/locate/cpc

Block iterative eigensolvers for sequences of correlated eigenvalue problems☆,☆☆

Edoardo Di Napoli [a,b,*], Mario Berljafa [c]

[a] Jülich Supercomputing Centre, Institute for Advanced Simulation, Forschungszentrum Jülich, Wilhelm-Johnen strasse, 52425 Jülich, Germany
[b] RWTH Aachen University, AICES, Schinkelstr. 2, 52062 Aachen, Germany
[c] Department of Mathematics, Faculty of Science, University of Zagreb, Bijenička cesta 30, 10000 Zagreb, Croatia

ARTICLE INFO

Article history:
Received 7 December 2012
Received in revised form
16 June 2013
Accepted 18 June 2013
Available online 26 June 2013

Keywords:
Density functional theory
Sequence of generalized eigenproblems
FLAPW
Block iterative eigensolvers
Eigensolver performance

ABSTRACT

In Density Functional Theory simulations based on the LAPW method, each self-consistent field cycle comprises dozens of large dense generalized eigenproblems. In contrast to real-space methods, eigenpairs solving for problems at distinct cycles have either been believed to be independent or at most very loosely connected. In a recent study (Di Napoli et al., 2012) [13], it was demonstrated that, contrary to belief, successive eigenproblems in a sequence are strongly correlated with one another. In particular, by monitoring the subspace angles between eigenvectors of successive eigenproblems, it was shown that these angles decrease noticeably after the first few iterations and become close to collinear. This last result suggests that we can manipulate the eigenvectors, solving for a specific eigenproblem in a sequence, as an approximate solution for the following eigenproblem. In this work we present results that are in line with this intuition. We provide numerical examples where opportunely selected block iterative eigensolvers benefit from the reuse of eigenvectors by achieving a substantial speed-up. The results presented will eventually open the way to a widespread use of block iterative eigensolvers in ab initio electronic structure codes based on the LAPW approach.

© 2013 Elsevier B.V. All rights reserved.

1. Introduction

Materials simulations based on Density Functional Theory [1] (DFT) methods have at their core a set of partial differential equations (Kohn–Sham [2]) which eventually lead to a non-linear generalized eigenvalue problem. Solving the latter directly is a daunting task and a numerical iterative self-consistent approach is preferred. It starts off by inputting an approximate electronic charge density to a cyclic process within which a linearized version of the eigenvalue problem is initialized and solved. At the end of each cycle a new charge density is computed and compared with the initial one. Self-consistency is reached when the distance between the input and output densities is below a certain required threshold; the process is then said to have converged. The entire simulation results in a series of so-called outer-iteration cycles often referred to as self-consistent field iterations.

Roughly speaking, all the existing DFT-based methods differ from each other by the choice of linearization scheme (also denoted as discretization), and by the choice of the effective Kohn–Sham (KS) potential. There are three discretization strategies commonly in use: (1) manipulation of localized functions (Gaussians, etc.), (2) expansion of the eigenfunctions over a plane wave basis set, and (3) discretization of the KS equations over a lattice in real space. While the first method is almost exclusively used in Quantum Chemistry the last two are widely used in Materials Science and present a series of pros and cons. The plane wave expansion leads to Hamiltonians with kinetic energy terms only on the main diagonal and are well suited to simulate solid crystals. In turn this discretization needs to approximate the Coulomb potential near the nuclei substituting it with a smooth pseudo-potential. In real-space discretization, potential terms in the Hamiltonian decay exponentially away from the diagonal [3] giving rise to quite sparse and large eigenvalue problems. This strategy is well suited mostly for disordered systems and insulators.

Among the plane wave strategies, the Full-Potential Linearized Augmented Plane Wave (FLAPW) [4,5] method constitutes one of the most precise frameworks for simulating transition metals and magnetic systems. The Kohn–Sham equations are discretized using a mix of radial and plane wave functions (see Section 2), parametrized by a vector $\mathbf{k}$ within the Brillouin zone of the momentum lattice. At each outer-iteration ℓ a set of eigenpencils $P_{\mathbf{k}}^{(\ell)}$,

☆ Article based on the research supported by the VolkswagenStiftung through the "Computational Science" fellowship.
☆☆ Preprints: AICES–2012/12–1 arXiv:1206.3768v1.
* Corresponding author at: Jülich Supercomputing Centre, Institute for Advanced Simulation, Forschungszentrum Jülich, Wilhelm-Johnen strasse, 52425 Jülich, Germany. Tel.: +49 0241 56001303.
E-mail addresses: e.di.napoli@fz-juelich.de, dinapoli@aices.rwth-aachen.de (E. Di Napoli), mberljaf@student.math.hr (M. Berljafa).

0010-4655/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.cpc.2013.06.017

labeled by **k**, is initialized and solved. Because FLAPW uses full-potential together with a partial plane wave expansion, each $P^{(\ell)}_{\mathbf{k}}$ is a dense and Hermitian generalized eigenvalue problem; its size n depends linearly – with a large pre-factor – on the number of atoms considered in a simulation, and typically ranges between 2000 and 20,000. Only a relatively small percentage of the bottom end of the spectrum is required, never exceeding 15%–20%, and often quite less.

In this work we consider sequences of generalized Hermitian eigenvalue problems as they arise in FLAPW. In this context a sequence is a set of N generalized eigenproblems identified by a progressive index ℓ

$$\{P^{(\ell)}\} \doteq P^{(1)}, \ldots, P^{(N)}; \qquad P^{(\ell)}: \quad A^{(\ell)}x = \lambda B^{(\ell)}x. \tag{1}$$

Within a sequence each eigenproblem is characterized by a Hermitian indefinite matrix A and a positive definite Hermitian matrix B. This setup is generally referred to as a matrix pencil or eigenpencil and it is known to have a bounded discrete spectrum with real positive and negative eigenvalues

$$\lambda_{\min} = \lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n = \lambda_{\max}. \tag{2}$$

Eigenpencils usually admit n distinct eigenvectors x_i satisfying a B-orthonormality relation $(x_i, Bx_j) = \delta_{ij}$ even when they correspond to identical eigenvalues. While in general $B \neq I$, in the special case $B = I$ the eigenpencil becomes a standard eigenvalue problem, and the orthonormality relation reduces to the standard $(x_i, x_j) = \delta_{ij}$.

In current codes implementing FLAPW [6–10], each sequence of eigenpencils $\{P^{(\ell)}\}$ is handled very much as a set of uncorrelated problems: each $P^{(\ell)}$ is solved in complete isolation from any other and independently passed as input to a prepackaged eigensolver of a standard library – like LAPACK [11] or its parallel version ScaLAPACK [12] – which outputs the desired portion of eigenspectrum and corresponding eigenvectors. The eigensolver is thus used as a black box and has no knowledge of the eigenproblems' spectral properties nor of the application from which they originated. As much as this process grants standardization and reliability, it is also far from being optimal. What is "lost in translation" is the possibility to render manifest the correlation between eigenpencils of the sequence $\{P^{(\ell)}\}$ in terms of precise numerical properties which are then passed to a solver that can exploit them.

In a recent work [13] it has been reported that eigenpencils with a successive outer-iteration index ℓ and the same **k**-vector are strongly correlated. Consequently, problems in a sequence are not only connected by a progressive index but, as for a sequence of numbers, there is a relation linking them. In FLAPW, such a numerical correlation become evident in the way the subspace angles between eigenvectors evolve from larger to smaller values as the sequence progresses toward higher outer-iteration indices [13]. It needs to be stressed that, contrary to what happens in real-space methods, the correlation between eigenvectors is a new and unexpected feature of FLAPW-based methods: since the eigenfunctions are delocalized and the function basis set is modified at each successive outer-iteration, it had been common belief that correlation was an unlikely phenomenon.

With evidence of the contrary in hand, it becomes natural to consider eigenvectors of $P^{(\ell)}$ as a set of approximate solutions that can be used by an appropriate eigensolver to accelerate the solution of $P^{(\ell+1)}$. The novelty of our contribution consists in showing that, by exploiting the collinearity between vectors of adjacent problems, we can significantly improve the performance of certain classes of eigensolvers. Since no eigensolver (QR, MRRR, Divide and Conquer, etc.) for dense problems accepts as input approximate eigenvectors, our strategy can only be carried out by using iterative eigensolvers. In the rest of the paper we first illustrate, through numerical experiments, the success of this strategy for three distinct block iterative eigensolvers, each representing a specific class of available methods. Then we focus on one of these solvers, develop a C language version and obtain similar results with emphasis on high-performance and scalability.

In Section 2 we first give a short description of how sequences of eigenproblems arise in DFT and how they translate into apparently uncorrelated dense eigenvalue problems. We then proceed to briefly report on the correlation between adjacent eigenproblems as illustrated in [13]. Our core results are presented in Section 3 where we introduce the selected block iterative eigensolvers followed by a description of the experimental setup and the numerical tests performed. We summarize our results in Section 4 and conclude with future work and acknowledgments.

2. Sequences of correlated eigenproblems

In this section we illustrate in some detail how sequences of eigenpencils arise in DFT. We start with a brief recall of the fundamentals of quantum mechanics, explain the need for an effective theory dealing with many particles and describe the FLAPW method self-consistent cycle. It is then shown why correlation among eigenproblems in a sequence is unexpected, and how the presence of such a correlation was exposed by looking at the evolution of eigenvectors as a function of the outer-iteration cycle index ℓ.

2.1. The rise of sequences in density functional theory

The electronic structure of a quantum mechanical system with L atoms and M electrons is described by the Schrödinger equation

$$H\Phi(x_1; s_1, \ldots, x_n; s_n) = \mathcal{E}\Phi(x_1; s_1, \ldots, x_n; s_n). \tag{3}$$

$H = -\frac{\hbar^2}{2m}\sum_{i=1}^{M}\nabla_i^2 - \sum_{i=1}^{M}\sum_{\mu=1}^{L}\frac{Z_\mu}{|x_i - a_\mu|} + \sum_{i<j}\frac{1}{|x_i - x_j|}$ is the Hamiltonian characterizing the dynamics of the electrons whose positions and spins are indicated by x and s respectively. $\mathcal{E}$ represents the energy of the system while Φ is the high-dimensional antisymmetric electronic wave function solving for Eq. (3). Already at this stage the Schrödinger equation looks very much like an eigenvalue problem, unfortunately one that is already very challenging to solve for values $M, L \geq 2$.

During the 1960s, a series of simplifications were introduced based on rigorous theorems [2,14] where the exact high-dimensional Eq. (3) was replaced by a large set of one-dimensional Kohn–Sham equations

$$\forall\, a \quad \text{solve} \quad \hat{H}_{\mathrm{KS}}\phi_a(\mathbf{r}) = \left(-\frac{\hbar^2}{2m}\nabla^2 + V_0(\mathbf{r})\right)\phi_a(\mathbf{r}) = \varepsilon_a\phi_a(\mathbf{r}). \tag{4}$$

The most important element in these equations is the substitution of the last two terms of H with an effective potential $V_0(\mathbf{r})[n]$ that functionally depends on the charge density $n(\mathbf{r})$: a function of all the one-particle wave functions $\phi_a(\mathbf{r})$. Because of this interdependence between V_0 and $\phi_a(\mathbf{r})$, Eq. (4) constitutes a set of non-linear partial differential equations.

Typically this set of equations is solved using an outer-iterative self-consistent cycle: it starts off with an initial charge density $n_{\mathrm{init}}(\mathbf{r})$, proceeds through a series of iterations and converges to a final density $n_N(\mathbf{r})$ such that $|n^{(N)} - n^{(N-1)}| < \eta$, with η as an a priori parameter. Convergence is achieved by an opportune mixing between output density $n_i(\mathbf{r})$ and one or more previous input densities $n_{\ell<i}(\mathbf{r})$. In the particular case of FLAPW the new

Download English Version:

https://daneshyari.com/en/article/10349527

Download Persian Version:

https://daneshyari.com/article/10349527

Daneshyari.com