

© 1999 Academic Press

Key Words: spectral projectors; density matrix; fast algorithms; wavelets; partitioned SVD.

1. INTRODUCTION

The goal of this paper is to introduce fast algorithms for computing spectral projectors. Although spectral projectors have a wide range of applications, we deal here primarily with density-matrix computations as they arise in the Kohn–Sham scheme.

The Kohn–Sham scheme provides a way to compute the ground state density of an arbitrary interacting system of n electrons. In a typical problem, given the positions and charges of nuclei, we would like to know the wave function describing the n electron ground state of the system. Since the wave function Ψ acts on $3n$ variables, it is too expensive to obtain. Instead, in density-functional theory (see, e.g., [1]) and other related theories one asks for the electron density, denoted $\rho_n(x)$, which gives the probability of finding an electron at x , i.e.,

$$\rho_n(x) = \int |\Psi(x, x_2, \dots, x_n)|^2 dx_2 \dots dx_n. \quad (1)$$

¹This research was partially supported by DARPA/NASA Grant S43 5-28646 (G.B, N.C., and M.J.M.) and DARPA/AFOSR Grant DOD F49620-97-1-0017 (G.B.).

In order to construct the density as the limit of an iterative process, the Kohn–Sham scheme produces a sequence of densities and potentials. At iteration i we compute a new (auxillary) potential from the previous density by

$$V_i(x) = V_0(x) + W[\rho_n^{i-1}](x), \quad (2)$$

where V_0 is the potential induced by the nuclei. Although the correct functional W for this process is not known, various approximations are in use (see, e.g., [1]).

To compute the new density, one finds the n smallest eigenvalues (with multiplicities) λ_j of the Hamiltonian $-\nabla^2 + V_i(x)$ and their corresponding eigenfunctions $\psi_j(x)$ and forms

$$\rho_n^i(x) = \sum_{j=1}^n |\psi_j(x)|^2. \quad (3)$$

If the densities $\rho_n^i(x)$ converge to some function $\tilde{\rho}_n(x)$, then the function $\tilde{\rho}_n(x)$ is the density of the system.

In order to avoid the costly computation of the eigenfunctions, the density can be constructed as the diagonal of the density matrix. Given a value μ such that λ

us to obtain linear scaling in the *number of atoms*. The efficiency of the linear scaling method depends on electron locality, but even when some electrons are poorly localized, most are well localized, so we can represent the non-local portion of the density matrix efficiently using singular value decompositions of appropriate pieces. We demonstrate our basic approach on a one-dimensional example and indicate considerations for two- and three- dimensional implementations.

For the case where the number of eigenfunctions included is large (in many other applications of spectral projectors) we present a multilevel partitioned representation of matrices (a technique due to Rokhlin and his collaborators [13–15]) which is based on singular value decompositions of submatrices. We explain the computational gain using the Christoffel–Darboux summation formula (see also [32]). We also present a method for partitioning the spectrum for the case where different sets of eigenfunctions require different bases for efficient representation.

Specifically, in Section 2 we define the matrix sign function and present the polynomial recursion to construct it. In Section 3 we develop tools to keep the matrices sparse during this recursion. We present the basic ideas within that section, and defer the details to Appendix A. In Section 4 we consider a numerical example to illustrate the claims of Section 3. In Section 5 we discuss extensions of these techniques to multiple dimensions and, finally, make concluding remarks in Section 6.

2. THE SIGN FUNCTION AND ASSOCIATED PROJECTORS

The ordinary sign function is defined on $(j, 1)$.

SP1 . 95701 . UT MOHLENKAMP
07

2.1. Recursive Construction

In this paper we use a polynomial recursion (also used in, e.g., [11, 12]) to compute $\text{sign}(T)$. The algorithm consists of the following steps:

$$\begin{aligned} T_0 & \text{D } T/kT^2 \\ T_{kC1} & \text{D } \frac{3T_k + T_k^3}{2}, \quad k \text{D } 0, 1, \dots \end{aligned} \quad (9)$$

Other polynomials may be used in place of the one above; see [12] for a discussion of the various choices.

We first demonstrate that $T_k \rightarrow \text{sign}(T)$ in (9). Observe that if U is the unitary transform that diagonalizes T_0 , then it also diagonalizes all T_k for $k \text{D } 1, 2, \dots$. Thus, we need only show that the scalar iteration $\lambda_{kC1} \text{D } (3\lambda_k + \lambda_k^3)/2$ converges to $\text{sign}(\lambda_0)$, provided that $|\lambda_0| \leq 1$. On the interval $[-1, 1]$ the function $(3\lambda + \lambda^3)/2$ is increasing and has the fixed points $\lambda \text{D } -1, 0, 1$ and no others. Since

$$\frac{\lambda_{kC1} - \lambda_k}{\lambda_k} \text{D } \frac{3}{2} + \frac{1}{2}\lambda_k^2, \quad (10)$$

we have either $0 < \lambda_k < \lambda_{kC1} \leq 1$ or $-1 \leq \lambda_{kC1} < \lambda_k < 0$. Therefore, $\lambda_k \rightarrow 0$ if $\lambda_0 \text{D } 0$, $\lambda_k \rightarrow 1$ if $0 < \lambda_0 \leq 1$, and $\lambda_k \rightarrow -1$ if $-1 \leq \lambda_0 < 0$. Both -1 and 1 are stable fixed points.

Remark 2.3. The iteration step in (9) is equivalent to the “purification transform” of McWeeny [17], where it is used as a correction in a variational method rather than as a recursion.

Remark 2.4. The following recursion (see, e.g., [10]) also provides an algorithm for computing $\text{sign}(T)$ for the matrix T ,

$$\begin{aligned} T_0 &\text{D} T \\ T_{k+1} &\text{D} T_k C T_k^{-1} \text{C}^{-1}, \quad k \text{D} 0, 1, \dots, \end{aligned} \quad (11)$$

where T_k^{-1} is a generalized inverse (if T has a null space). We avoid this formulation because it requires computing an inverse, and does not preserve the ordering of the eigenvalues. We also point to Appendix B where the spectral projector is expressed as an integral of Green’s function.

3. TOOLS FOR SPARSE REPRESENTATION

In this section we develop several representations for fast matrix–matrix multiplications within the recursion (9). We present these ideas briefly in this section, and defer estimates and proofs to Appendix A. We consider only one-dimensional problems here and mention considerations for multiple dimensions in Section 5.2.

The representations of this section are critical to our approach since they control the speed of the algorithm. We describe an adapted discretization of the Hamiltonian for a single atom in Section 3.1. In Section 3.2 we consider the sparsity of the spectral projector for several atoms and introduce additional structure into the representation. In Section 3.3 we demonstrate a method suitable for projectors where n , the number of eigenfunctions of interest, is large (e.g., above the Fermi level).

3.1. The Adapted Representation of the Hamiltonian

In order to construct the spectral projector using the sign function iteration (9) we must first convert the true Hamiltonian $H \text{D} \int_{\mathbb{R}^2} \text{C} V(x)$ to matrix form. This can be done by either sampling in space or representing the operator in some basis. We will represent it on a basis because this will allow us access to both the space and the spatial-frequency domains. We will consider only orthonormal bases. We have observed that (quite naturally) the way in which the initial discretization is handled has a strong effect.

The representation of H on some finite set of basis functions can be viewed as a projection of H onto a subspace. We will call this projector P a “rough projector” and apply the iteration (9) to the matrix $\tilde{H} \text{D} PHP$. The projector P identifies the subspace spanned by the first n eigenfunctions. In discretizing the original Hamiltonian, we would like to project it on a subspace that (i) includes the subspace indicated by P with controlled accuracy, and (ii) is not significantly larger than the subspace indicated by P . Formally, these conditions mean $\|PP^\dagger - P\| < \epsilon$ for some desired accuracy ϵ , and the operator $P(I - P)$ has small rank.

In order to construct the rough projector P we need basis functions with controlled localization in both space and frequency. The necessity of this localization is implied by an

Note that P does not need to “resolve” or “capture” H itself, but only P . Other eigenvalues and eigenvectors of H will be changed by P , but an orthogonal projection does not change the *sign* of the eigenvalues. See Section A.1 for a more detailed discussion of the effect of a rough projector.

We choose P to be a projection onto a collection of wavelets. The design of P is based on the potential $V(x)$ and eigenvalue bound μ only. The size of $\mu \ll V(x)$ determines the maximal “instantaneous frequency” and therefore the necessary sampling rate (i.e., wavelet subspace). The derivative of $V(x)$ determines how much “frequency spillage” we will have, and thus how well P can match P locally. We will characterize this subspace using an instantaneous frequency perspective. In Section A.2 we provide a rigorous justification using the local cosine basis (see, e.g., [18, 19]) and we indicate below how to translate this to the wavelet basis.

The eigenfunction $\psi_n(x)$ satisfies (by definition) $\psi_n''(x) + (\lambda_n - V(x))\psi_n(x) = 0$. The WKB (quasi-classical) approximation predicts behavior like

$$\exp\left(\pm i \int_{x_0}^x \sqrt{\lambda_n - V(t)} dt\right) \quad (12)$$

and thus instantaneous frequency $\sqrt{\lambda_n - V(x)}$. Intuitively this says that ψ_n “lives” on the curve $\xi = \nu_n(x) = \sqrt{\lambda_n - V(x)}$ in the $x \in \xi$ (space \in spatial-frequency) phase plane. On the phase plane a local cosine basis element is viewed as a rectangle with x -support on its base interval, shifted in ξ by its frequency, with area a constant (depending on the normalization). Intuitively, those boxes that intersect $\nu(x)$ should correspond to local cosine elements that yield significant coefficients (see Fig. 1).

The important conclusion from the estimates in Section A.2 is that for potentials of the form $V(x) = C/x^2$ (in one dimension), the number of local cosine basis functions needed to construct P is proportional to $\sqrt{C} \dots n$, where n is the rank of P . We can thus represent




FIG. 1. Schematic of instantaneous frequency plots for several ψ_k with potential $V(x) = C/x^2$, overlaid with the local cosine subdivision for the subspace used for P .




FIG. 2. Schematic of instantaneous frequency plots for ψ_k overlaid with a wavelet phase plane.

\tilde{H} as an $O(n \times n)$ matrix in this adapted local cosine coordinate system and compute P in $O(n^3)$ time using the recursion (9). For a single atom this result may already be sufficient for the fast computation of P , since n is never very large. In the following sections we give further representation techniques to deal with multiple atoms and the case where the number of eigenfunctions is large.

To translate the above results to a wavelet representation, we need only note that the wavelet partition of the phase plane is compatible with the type of partition desired for P . In particular, high frequency is associated with small spatial support and high change in frequency (see Fig. 2). We therefore can represent \tilde{H} and P as $O(n \times n)$ matrices in an adapted wavelet subspace. The constant involved will depend on the choice of wavelet and the desired precision ϵ . The dependence on ϵ for a wavelet expansion is generally $\log(1/\epsilon)$, yielding matrices of size $O(n \log(1/\epsilon) \times n \log(1/\epsilon))$.

In what follows we use the standard form of the matrices, which is equivalent to simply changing our system of coordinates into the wavelet basis, and note that it is also possible to use the non-standard form of [20, 21].

Remark 3.1. One could construct P using the atomic orbitals. At low precision this should perform well, since the atomic orbitals match the eigenfunctions well. At higher precision, however, atomic orbitals are a poor choice because they do not allow local refinements adapted to the particular potential in use.

Remark 3.2. If we choose the subspace for P “too small” the density constructed will still be an approximation of the true density (see Section A.1). Early in the Kohn–Sham iteration it may even be desirable to use an approximate density.

3.2. Multiple Atoms

In this section we consider the case where we have N nuclei, each with n electrons. Besides the rough projection P , we will need additional structures for the sparse representation

of our matrices. We design our representation so that the final answer, the projector P , will be sparse. The initial Hamiltonian is banded and so will also be sparse in our representation. It appears that due to the monotonicity properties of the recursion (9), choosing a representation so that the initial Hamiltonian and final projector are sparse is sufficient for maintaining the sparsity of the intermediate matrices. In this section we first produce linear scaling in N , as others have done. Second we provide additional structures to reduce the constant in the complexity estimates.

As a first step we note that at any point in the recursion many of the matrix entries may be less than ϵ and can be neglected, leaving a sparse matrix. For a single atom and moderate number of eigenfunctions n , such a sparse representation may provide only a limited benefit. However, in the case where we have N nuclei, each with n electrons, this sparse representation provides a significant benefit. Without it we would have $O(nN \times nN)$ matrices and therefore an $O((nN)^3)$ algorithm.

As others (e.g., [2, 3, 22–24]) have noted, the locality of the electrons (eigenfunctions) allows us to obtain linear scaling in the number of nuclei N . The basic argument is as follows: For an eigenfunction to represent a bound state electron, it must be L^2 normalized and hence decay at infinity. Methods like WKB allow us to estimate the decay of the eigenfunction and thus the region where it is numerically significant. Suppose the amplitude of the eigenfunctions is significant for d units, where atoms are s units apart and $d > s$. If

operations. The constant has been significantly reduced, but may still be too large for some problems. The fact that wavelets are well suited to representing wave functions has been noted in [25].

If the number of nuclei in interaction range (d/s) is large, we will need an additional technique. The block B_{ij} is formed by eigenfunctions that have a significant component near both nuclei i and j . The number of entries in B_{ij} is determined by the highest energy eigenfunction ψ_n . This eigenfunction has the slowest rate of decay, and so we expect the off-diagonal blocks to remain full ($O(n \times n)$), and decay (perhaps slowly) in amplitude as $|i - j|$ increases. Lower energy eigenfunctions, however, will decay much more rapidly, so B_{ij} , although full, will become *low rank* (up to ϵ) long before $|i - j| > d/s$. We can represent these blocks using the singular value decomposition (SVD) and obtain a much more efficient representation. This technique takes advantage of the fact that core electrons interact only at short distances.

3.3. *Partitioned SVD Representations*

As the number of eigenfunctions n increases, the cost of computation using wavelet compression may increase like n^3 . In physical systems the number of eigenfunctions per atom is never very large, and the localization of the eigenfunctions keeps the representation sparse.




FIG. 3. Partitioned SVD.

Dividing this by $a_{nC_1}(x_j - y)$ reveals the top term in the sum, and we are left to prove the same theorem with n replaced by $n - 1$. The theorem follows by induction. ■

orthogonal polynomials by changes of variables. In particular it holds for the Chebyshev polynomials $T_n(x)$, which under the change of variables $x = \cos \theta$ become $T_n(\cos \theta) = \cos(n\theta)$

4. NUMERICAL EXAMPLES

In this section we test our methods on a simple one-dimensional example. The first question we wish to answer is how many iterations are required for convergence, as a function of the condition number of the initial matrix. Second we test how well our wavelet and PSVD representations compare to direct (sparse) matrix multiplications. This comparison is done for both high and low μ (equivalently n).

To determine the number of iterations required we consider the operator T , a 512×512 discretization of ∂^2 with periodic boundary conditions. Since the smallest eigenvalue of T is 0, $T_0 = T/\|T\|_C \kappa^{-1} I$ has condition number κ . (In the general problem, the condition number is the band gap divided by the matrix norm.) We then see how many iterations it takes for T_0 to converge to the identity to accuracy ϵ for several values of κ . According to Section 2.1, we should have convergence in $O(\log \kappa + \log \log(1/\epsilon))$ iterations, and the $O(\log \kappa)$ portion is confirmed in Table I. See [12] for a more detailed analysis of the convergence of this iteration.

To test the sparsity of our representations we consider a 512×512 finite-difference discretization of the operator ∂^2 on the interval $[1/2, 1/2]$ with periodic boundary conditions. In our first example, we choose $\mu \gg 0$ and construct the projector onto the 15 eigenfunctions with eigenvalues less than μ . In our second example, we choose μ so that we project onto 70 eigenfunctions. In all cases we use a fifth-order finite difference




FIG. 6. (Example 4.1) Relative error in the density generated in Fig. 5.




FIG. 7. (Example 4.2) Ratio of the number of significant coefficients above the threshold 10^{-8} to the total number of matrix elements during the recursion (9) for μ capturing 70 eigenvectors.




FIG. 8. (Example 4.2) Relative error in the density generated in Fig. 7.

the PSVD still performs well. In Fig. 8 we plot the error in the density for this example. We chose μ between two eigenvalues that are paired, so the band gap is small and $\kappa \dots 10^6$. We achieve slightly better performance than in Fig. 6 simply because we measure relative error.

5. EXTENSIONS

5.1. *Partitioning the Spectrum*

The case where $\mu > 0$ can be treated by applying an appropriate shift and then using the construction above.

Such partitioning is useful if for some reason computing $P^0(x, y)$ and then $P^{00}(x, y)$ is easier than simply computing $P_n(x, y)$. Within the Kohn–Sham scheme, the lower parts of the spectrum, which correspond to densities which are localized around the nucleus, should converge more quickly under the DFT-LDA iteration than other parts of the spectrum. We may then be able to fix P^0 early in the iteration process and save some work. This idea is similar to the use of pseudopotentials except that we have not modified the potential, but instead the entire operator. We note that H^0 is no longer of the form “Laplacian plus potential.”

5.2. Multidimensional Implementations

Efficient implementation of both wavelet and PSVD representations in multiple dimensions requires careful attention. The straightforward generalization, although available, is not efficient.

In Example 4.1 we demonstrated that choosing a good initial adapted representation was crucial for efficiency. In multidimensional problems the treatment of the singularities (e.g., the Coulomb potential of the ions) will also become critical. We plan to use multiwavelets as a tool of discretization in multiple dimensions. These bases allow us to position boxes so that the point singularities of the ionic potential coincide with the corners of the parallelograms where the multiwavelets are supported. At these corners the multiwavelets are discontinuous already and so should be able to match the singularity with fewer scales than any overlapping wavelet basis. A paper on this topic which is a follow-up to [27, 28] is in preparation [29].

In addition, in a separate work [30] it is shown that for a large class of operators the difference between the operator and its projection on a coarse scale can be represented as a (small) sum of separable operators. This approach is shown to produce an efficient generalization for multidimensional implementation in, e.g., wavelet bases. We plan to use these results as a way of implementing the constructions of this paper in multiple dimensions.

6. CONCLUSIONS

We have presented a fast algorithm for the construction of a spectral projector. This algorithm allows us to compute the density matrix, as used in, e.g., the Kohn–Sham iteration, and so obtain the electron density. We computed the spectral projector by constructing the matrix sign function through a simple polynomial recursion. We have presented several techniques for fast computation within this recursion, using bases with controlled space–spatial frequency localization.

Since spectral projectors appear in many contexts, we expect many additional applications of our approach. In particular we expect this basic approach to work in molecular dynamics simulations and homogenized wave propagation. We note that the details of the appropriate representation to maintain sparsity may vary.

APPENDIX A: MATHEMATICAL ESTIMATES

A.1. The Effect of a “Rough” Projection

In this section we examine the effect of the rough projector from Section 3.1. We show that the signs of the eigenvalues are preserved. This fact is closely related to the law of

THEOREM A.1. *Let I be an interval of length l , ψ_i the normalized eigenfunction as*

fails only if $(jV^{(k)}(x))^{1/(kC_1)}/k$ increases as a function of k , which is not the case in our example.

One conclusion we may draw is that the sampling rate remains finite as we approach the singularity, as long as we only wish to capture ψ_n up to ϵ . The second, more important, conclusion is that \mathcal{P}

A.3. Spectral Equivalence

For the techniques in Section 3.3 to be valid in our case, we need estimates to show that $\int_{\mathbb{R}^2} |\hat{f}(x)|^2 dx$ and $\int_{\mathbb{R}^2} |\hat{f}(x)|^2 C V(x) dx$ are “spectrally equivalent.” In this section we give estimates showing to what degree this is true. Interpreting these results as good or bad will depend on the particular situation.

Theorem A.1 gives decay of local cosine coefficients like $j!^2 V'_{j+1} (p\pi)^{2j+1}$. In Section A.2 we used this for small V' and large p , but it can also be used for large V' and small p . It is slightly more convenient to compute the length of the set f




FIG. 9. The contour C , composed of C_I and C_{II} .

where the contour C is shown in Fig. 9, T is a self-adjoint matrix, and at no point on the contour C is the matrix $T - zI$ singular.

We can write the integral (37) as

$$\int_C \frac{1}{2\pi i} (T - zI)^{-1} dz = \int_C \frac{1}{2\pi i} U^{-1} (D - zI)^{-1} U dz, \quad (38)$$

where U is a unitary transformation which diagonalizes T and D is a diagonal matrix such that $T = U^{-1} D U$. The integral (37) may then be evaluated element-by-element on D , as in

$$\int_C \frac{1}{2\pi i} (\lambda - z)^{-1} dz, \quad (39)$$

where λ is a diagonal element of D .

We define the parts of the contour C as in Fig. 9. The vertical part, C_I , runs from $\mu + iM$ to $\mu - iM$. The part labeled C_{II} is a circular arc running from $\mu - iM$ to $\mu + iM$.

16. C. Bischof, X. Sun, A. Tsao, and T. Turnbull, A study of the invariant subspace decomposition algorithm for banded symmetric matrices, in *Proceedings, Fifth SIAM Conference on Applied Linear Algebra, 1994*.
17. R. McWeeny, Some recent advances in density matrix theory, *Rev. Mod. Phys.* **32**(2), 335–369 (1960).
18. R. R. Coifman and Y. Meyer, Remarques sur l'analyse de Fourier à fenêtre, *C.R. Acad. Sci.* **312**(1), 259–261 (1991).
19. M. V. Wickerhauser, *Adapted Wavelet Analysis from Theory to Software* (A. K. Peters, Boston, 1994).
20. G. Beylkin, R. Coifman, and V. Rokhlin, Fast wavelet transforms and numerical algorithms i, *Commun. Pure Appl. Math.* **44**, 141–183 (1991).
21. D. L. Gines, G. Beylkin, and J. Dunn, LU factorization of non-standard forms and direct multiresolution solvers, *Appl. Comput. Harmon. Anal.* **5**, 156–201 (1998). [PAM Report 278, 1996]
22. P. Ordejón, D. A. Drabold, R. M. Martin, and M. P. Grumbach, Linear system-size scaling methods for electronic-structure calculations, *Phys. Rev. B* **51**(3), 1456–1476 (1995).
23. P. Ordejón, E. Artacho, and J. M. Soler, Self-consistent order- N density-functional calculations for very large systems, *Phys. Rev. B* **53**(16), 10441–10444 (1996).
24. E. Hernández and M. J. Gillan, Self-consistent first-principles technique with linear scaling, *Phys. Rev. B* **51**(15), 10157–10160 (1995).