Journal of Computational Physics **152,** 32–54 (1999)
Article ID jcph.1999.6215, available online at http://www.idealibrary.com on IDEAL®

Fast Spectral Projection Algorithms for Density-Matrix Computations[1]

Gregory Beylkin, Nicholas Coult, and Martin J. Mohlenkamp

Department of Applied Mathematics, University of Colorado at Boulder, Boulder, Colorado 80309-0526
E-mail: beylkin@paley.colorado.edu, coult@littlewood.colorado.edu, mjm@colorado.edu

Received August 17, 1998; revised January 20, 1999

We present 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 compute the spectral projector by constructing the matrix sign function through a simple polynomial recursion. We present several matrix representations for fast computation within this recursion, using bases with controlled space–spatial-frequency localization. In particular we consider wavelet and local cosine bases. Since spectral projectors appear in many contexts, we expect many additional applications of our approach. © 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.,

$$\rho_n(x) = \int |\Psi(x, x_2, \ldots, x_n)|^2 \, dx_2 \cdots dx_n. \tag{1}$$

[1] This 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

$$V_i(x) = V_0(x) + W\left(\rho_n^{i-1}(\cdot), x\right), \tag{2}$$

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) $\{\lambda_j\}$ of the Hamiltonian $-\nabla^2 + V_i(x)$ and their corresponding eigenfunctions $\{\psi_j(x)\}$ and forms

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

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 $\lambda_n < \mu < \lambda_{n+1}$, the density matrix is defined as

$$P_\mu(x, y) = P_n(x, y) = \sum_{\lambda_j < \mu} \psi_j(x)\bar{\psi}_j(y). \tag{4}$$

Although we have defined it using the eigenfunctions, the density matrix can be constructed without computing the eigenfunctions explicitly. The density-matrix approach is used in several methods. For example, in [2, 3] the density matrix is constructed using a variational approach. In [4–8], the authors use the Chebyshev polynomials to approximate the density matrix. Although not directly related to the computation of the spectral projector and density matrix, we would like to mention [9] and the references therein, where fast methods for the computation of the density of states are developed.

In our approach we compute the density matrix as a spectral projector by computing the sign function. The standard methods for computing the sign function may be found in the survey [10]. More recently in [11, 12], polynomial recursions for the sign function are applied to the problem of computation of eigensystems, with the goal of parallelizing such computations. The attractive feature of polynomial recursions versus Chebychev polynomial approximations is that the order of approximation grows exponentially rather than linearly. The main difficulty in using polynomial recursions is that they require matrix–matrix multiplications and thus are not suitable for large-size problems. The key point of our approach is that we consider matrix representations that remain sparse (up to finite but arbitrary accuracy) throughout the iteration that produces the spectral projector.

We present a set of tools for the fast computation of the sign function within the polynomial iteration. We observe that it is very important to construct an efficient representation of the Hamiltonian, which is the starting point for our iteration. To this end we use a “rough” projection of the operator onto an adapted wavelet subspace sufficient to represent the density matrix. We show that the wavelet system of coordinates provides a suitable choice for projectors corresponding to the lowest eigenvalues. Due to physical considerations, the computational cost will scale cubicly in the number of electrons *per atom* (which is never a very large number). Using a sparse representation within this rough subspace allows 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 conluding remarks in Section 6.

2. THE SIGN FUNCTION AND ASSOCIATED PROJECTORS

The ordinary sign function is defined on $(-\infty, \infty)$ by

$$\operatorname{sign}(\lambda) = \begin{cases} 1, & \lambda > 0 \\ 0, & \lambda = 0 \\ -1, & \lambda < 0. \end{cases} \tag{5}$$

For a matrix or operator of the form

$$T(x, y) = \sum_j \lambda_j \psi_j(x) \bar{\psi}_j(y) \tag{6}$$

with λ_j real and $\{\psi_j\}$ an orthonormal set, we define

$$\operatorname{sign}(T)(x, y) = \sum_j \operatorname{sign}(\lambda_j) \psi_j(x) \bar{\psi}_j(y). \tag{7}$$

We construct the spectral projector from the sign function using

$$P_\mu(x, y) = \sum_{\lambda_j < \mu} \psi_j(x) \bar{\psi}_j(y) = (I - \operatorname{sign}(T - \mu I))/2. \tag{8}$$

Remark 2.1. For non-self-adjoint T, the sign function in (5) is defined as the sign of the real part of λ.

Remark 2.2. We could use the Heaviside function $(I - \operatorname{sign}(-x))/2$ instead of the sign function in (7) and construct the projector in the analogous way.

2.1. *Recursive Construction*

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

$$\begin{aligned} T_0 &= T/\|T\|_2 \\ T_{k+1} &= \left(3T_k - T_k^3\right)/2, \qquad k = 0, 1, \dots. \end{aligned} \tag{9}$$

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

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

$$\frac{\lambda_{k+1}}{\lambda_k} = \frac{3}{2} - \frac{1}{2}\lambda_k^2 \ge 1, \tag{10}$$

we have either $0 < \lambda_k < \lambda_{k+1} \le 1$ or $-1 \le \lambda_{k+1} < \lambda_k < 0$. Therefore, $\lambda_k \equiv 0$ if $\lambda_0 = 0$, $\lambda_k \to 1$ if $0 < \lambda_0 \le 1$, and $\lambda_k \to -1$ if $-1 \le \lambda_k < 0$. Both -1 and 1 are stable fixed points.

The number of iterations needed for (9) to converge to accuracy ϵ is $\mathcal{O}(c \log_2 \kappa + \log_2 \log_2(1/\epsilon))$, where κ is the condition number of T_0. At the beginning of the iteration, the smallest (say positive) eigenvalue of T_0 is like κ^{-1}. Expanded about $\lambda = 0$, the scalar recursion is approximately $\lambda_{k+1} = (3/2)\lambda_k$. The number of iterations needed for κ^{-1} to reach some intermediate value like $1/2$ is thus $\mathcal{O}(\log_{(3/2)} \kappa)$. Now we enter the regime where $(1 - \lambda_k)$ is small and convergence is quadratic. From this point to achieve precision ϵ takes $\mathcal{O}(\log_2 \log_2(1/\epsilon))$ iterations.

The recursion (9) requires matrix–matrix multiplications and so still appears to be costly as a computational tool. We, however, provide a mechanism for maintaining sparsity of the matrices during the recursion. In [16], symmetric band reductions are used for this purpose. Our approach is based on using wavelets or other representations (local cosine, multiwavelets, partitioned SVD) to maintain sparsity during the recursion. In the future we plan to investigate whether the approach of [16] can be incorporated into our approach as well.

One of the advantages of (9) is demonstrated in the proof above, namely, a monotone evolution of eigenvalues throughout the iteration. We observe that selecting a representation so that the original Hamiltonian and the resulting sign function are sparse appears to be sufficient for maintaining sparsity in the intermediate matrices.

One simple but very important observation is that the compressibility (sparse representation for finite but arbitrary accuracy) of projectors may be different than that of the individual eigenfunctions. We also note that computation of the projections via (9) has qualitatively different properties than that of the direct computation of individual eigenfunctions. Namely, if an eigenvalue has high multiplicity, the algorithms for direct computation of the eigenfunctions have difficulties, whereas the only difficulty that might occur in (9) is a possible loss of accuracy if T_0 has a nullspace or a high condition number.

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 $\operatorname{sign}(T)$ for the matrix T,

$$\begin{aligned} T_0 &= T \\ T_{k+1} &= \left(T_k + T_k^{-1}\right)/2, \qquad k = 0, 1, \ldots, \end{aligned} \tag{11}$$

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 $H = -\nabla^2 + V(x)$ 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 $\mathcal{P}$ a "rough projector" and apply the iteration (9) to the matrix $\tilde{H} = \mathcal{P}H\mathcal{P}$. The projector P identifies the subspace spanned by the first n eigenfunctions. In discretizing the origional 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 $\|\mathcal{P}P - P\| < \epsilon$ for some desired accuracy ϵ, and the operator $\mathcal{P}(I - P)$ has small rank.

In order to construct the rough projector $\mathcal{P}$ we need basis functions with controlled localization in both space and frequency. The necessity of this localization is implied by an analysis of the instantaneous frequency of the eigenfunctions. It is clear that to be efficient, $\mathcal{P}$ must be adapted to the potential $V(x)$ and the energy cutoff μ. Near the nucleus the eigenfunctions are more oscillatory, so $\mathcal{P}$ must allow higher frequencies.

Note that $\mathcal{P}$ does not need to "resolve" or "capture" H itself, but only P. Other eigenvalues and eigenvectors of H will be changed by $\mathcal{P}$, 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 $\mathcal{P}$ to be a projection onto a collection of wavelets. The design of $\mathcal{P}$ is based on the potential $V(x)$ and eigenvalue bound μ only. The size of $\mu - V(x)$ 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 $\mathcal{P}$ 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) $\psi_n''(x) = -(\lambda_n - V(x))\psi_n(x)$. The WKB (quasi-classical) approximation predicts behavior like

$$\exp\left(\pm i \int^x \sqrt{\lambda_n - V(t)}\, dt\right) \tag{12}$$

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 $x \times \xi$ (space × 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 $-C/|x|$ (in one dimension), the number of local cosine basis functions needed to construct $\mathcal{P}$ is proportional to $\sqrt{C} \approx n$, where n is the rank of P. We can thus represent

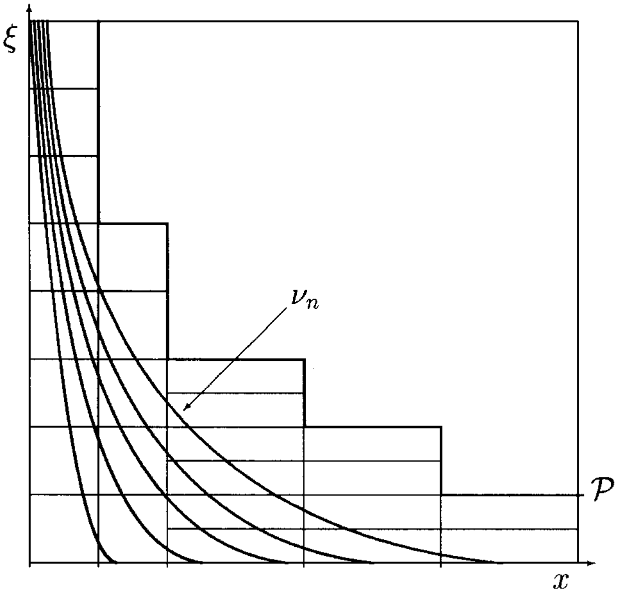

FIG. 1. Schematic of instantaneous frequency plots for several ψ_k with potential $V(x) = -C/x$, overlaid with the local cosine subdivision for the subspace used for $\mathcal{P}$.

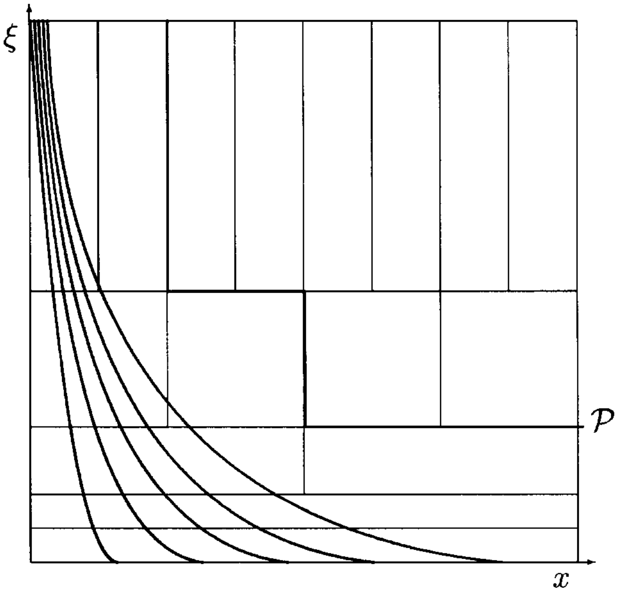

FIG. 2. Schematic of instantaneous frequency plots for ψ_k overlaid with a wavelet phase plane.

$\tilde{H}$ as an $\mathcal{O}(n \times n)$ matrix in this adapted local cosine coordinate system and compute P in $\mathcal{O}(n^3)$ 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 $\mathcal{P}$. 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 $\mathcal{O}(n \times n)$ 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 $\mathcal{O}(n \log(1/\epsilon) \times n \log(1/\epsilon))$.

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 $\mathcal{P}$ 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 $\mathcal{P}$ “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 $\mathcal{P}$, 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 $\mathcal{O}(nN \times nN)$ matrices and therefore an $\mathcal{O}((nN)^3)$ 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 we use a K point discretization per unit, then the matrix is banded, with bandwidth $d \cdot K$ and diagonal length $N \cdot s \cdot K$. It thus has $N \cdot s \cdot d \cdot K^2$ entries and can be multiplied in the recursion in $N \cdot s \cdot K \cdot (d \cdot K)^2$ operations. We have obtained linear scaling in N, but the constant $s \cdot d \cdot K^2$ (or $s \cdot K \cdot (d \cdot K)^2$) may be quite large.

By using the wavelet coordinate system and the rough projector we can reduce this constant. The matrix can be organized into a set of blocks, each representing the interaction between a pair of nuclei (including self-interactions). We will discuss this structure here at a fixed (wavelet) scale and act as if there were only this scale and a coarse scale. Since the main estimate used for the blocksize is $\mathcal{O}(n \times n)$, which includes the contributions from all scales, in what follows we need not sum over the scales to obtain the overall estimate. We decompose the projector into matrix blocks B_{ij}, $1 \leq i, j \leq N$. Letting ψ_k^a denote eigenfunction k of atom a,

$$B_{ij} = \sum_{a=1}^{N} \sum_{k=1}^{n} \hat{\psi}_k^a(x)\big|_{x \approx i} \overline{\hat{\psi}_k^a}(y)\big|_{y \approx j}, \tag{13}$$

where $\hat{}$ denotes the wavelet transform, x and y are wavelet coordinates, and the restriction $x \approx i$ means that x is near nucleus i. Each block is an $\mathcal{O}(n \times n)$ square of (possibly) significant entries, surrounded by entries less than ϵ. This empty area between blocks is the classically forbidden region. In this region ψ_k is a smooth (decaying) function, which can be represented by a small number of coefficients at a coarse scale. The rough projector $\mathcal{P}$ removes from the Hamiltonian those coordinates that are not needed for the projector, thus deleting the empty space between blocks. For an eigenfunction to contribute to B_{ij} it must have a significant component near both nuclei i and j. The assumptions that all eigenfunctions are only significant for d units and nuclei are s units apart mean that B_{ij} is set to zero if $|i - j| > d/s$. Our matrix thus is block-banded, with N blocks along the diagonal, d/s blocks on each row, and with blocks of size $\mathcal{O}(n \times n)$. Thus the matrix has $\mathcal{O}(N \cdot (d/s) \cdot n^2)$ entries and can be multiplied in the recursion in $\mathcal{O}(N \cdot (d/s)^2 \cdot n^3)$

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 $(\mathcal{O}(n \times n))$, 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 eigenfuctions 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. In other applications of projectors, we may have a situation corresponding to keeping a large number of eigenfunctions (above the Fermi level) on a single atom. For this case, we propose a technique that should be insensitive to n, or even improve for large n. The representation we use was developed by Rokhlin and his collaborators [13–15]. We use it to exploit the implications of the Christoffel–Darboux summation formula (see [32]). This theorem does not apply in our case, but a similar approximate result does hold. We first present the Christoffel–Darboux summation formula and a sketch of Rokhlin's approach. Then we present the approximate result that holds in our context.

PROPOSITION 3.3. *Let $\{p_k(x)\}_{k=0}^{\infty}$ be a set of normalized orthogonal polynomials on $[-1, 1]$ constructed from $1, x, x^2, \ldots$ by Gramm–Schmidt orthonormalization with respect to some weight $w(x)\,dx$. Then $\{p_k(x)\}_{k=0}^{\infty}$ satisfies a three-term recurrence initialized by $p_{-1}(x) = 0,\ p_0(x) = a_0$, and with general term $p_n(x) = (a_n x + b_n) p_{n-1}(x) + (a_n/a_{n-1}) p_{n-2}(x)$.*

THEOREM 3.4 (Christoffel–Darboux [26, p. 43]).

$$\sum_{k=0}^{n} p_k(x) p_k(y) = a_{n+1}^{-1} \frac{p_{n+1}(x) p_n(y) - p_n(x) p_{n+1}(y)}{x - y}. \tag{14}$$

Since the proof is brief, we present it here.

Proof. Expanding the numerator using the recurrence relation we obtain

$$\begin{aligned} &\left((a_{n+1}x + b_{n+1}) p_n(x) + \frac{a_{n+1}}{a_n} p_{n-1}(x)\right) p_n(y) \\ &\qquad - p_n(x)\left((a_{n+1}y + b_{n+1}) p_n(y) + \frac{a_{n+1}}{a_n} p_{n-1}(y)\right) \\ &= a_{n+1}(x - y) p_n(x) p_n(y) + (a_{n+1}/a_n)(p_n(x) p_{n-1}(y) - p_{n-1}(x) p_n(y)). \end{aligned} \tag{15}$$

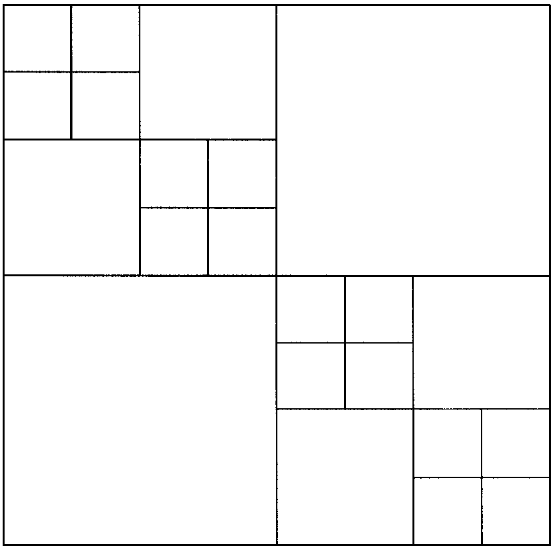

FIG. 3. Partitioned SVD.

Dividing this by $a_{n+1}(x-y)$ 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. ■

This theorem says that the projector onto the first n polynomials is significantly less complicated than the set $\{p_k(x)\}_{k=0}^{n}$. The kernel $(p_{n+1}(x)p_n(y))/(x-y)$ naturally factors into

$$p_{n+1}(x)\frac{1}{x-y}p_n(y) = D_1 A D_2. \tag{16}$$

The matrices D_1, D_2 are diagonal and can be applied with ease. A is not diagonal, but is compressible when partitioned as in Fig. 3. On each off-diagonal square A is a low-rank matrix in the sense that its singular value decomposition has only a few singular values above a given threshold. D_1AD_2 must then also be of low rank (verified by direct evaluation). To represent D_1AD_2 on a square, we need to store some number of singular values, and the corresponding singular vectors, each with length the size of the current square. The cost to represent the entire matrix is the same as for A in this representation, which is $\mathcal{O}(K(\log K)^2)$ for a K point discretization. We call this a partitioned SVD representation (PSVD).

Remark 3.5. It might be possible to compress the singular vectors as well. The low rank of A is not sufficient to imply that D_1AD_2 is compressible in a wavelet basis, however. For compressiblity of D_1AD_2 we would also need the vectors which are the diagonals of D_1 and D_2 to be compressible. For example, if D_1 and D_2 are random diagonal matrices, the product D_1AD_2 appears random to the wavelets, and no compression is possible.

When matrices in PSVD form are multiplied, it creates multiple contributions to each part of the output matrix. To return to the PSVD form, we must add these contributions and move them to the correct size square. This process is similar to multiplying in non-standard wavelet form [21].

The projectors we are considering are not formed from orthogonal polynomials and so Theorem 3.4 does not directly apply. The theorem does hold for objects derived from 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 $T_n(\cos\theta) = \cos(n\theta)$. Theorem 3.4 thus holds for cosines, after the appropriate change of variables. The cosines are eigenfunctions for the Hamiltonian $-\nabla^2$ which we will consider as our prototype, with other Hamiltonians a perturbation of this by a potential.

We note that for the eigenfunctions of $-\nabla^2$ we could bypass Christoffel–Darboux and use instead the formula for the projector onto exponentials (valid in any dimension),

$$\sum_{k=0}^{n} e^{ikx}e^{-iky} = \frac{1-e^{i(n+1)(x-y)}}{1-e^{i(x-y)}}. \tag{17}$$

The approximate version of Theorem 3.4 that holds in our case is based on the claim that $-\nabla^2$ and $-\nabla^2 + V(x)$ are "spectrally equivalent" in the following sense: Let R_m be the projector onto the first m cosines (or exponentials). We will decompose P as

$$P = R_m - (I-P)R_m + (I-R_m)P \tag{18}$$

(see Fig. 4). The Christoffel–Darboux theorem applies to R_m, so it can be represented efficiently in the PSVD. We claim that there is an m so that $(I-R_m)P$ and $(I-P)R_m$ are either low rank or highly localized. Supposing that $(I-R_m)P$ or $(I-P)R_m$ is of rank r, it adds at most r singular values to each PSVD square, and thus at most $rK\log K$ additional coefficients in the PSVD representation. If $(I-R_m)P$ or $(I-P)R_m$ were highly localized, it would interact with only a few squares, and thus add few additional coefficients in the PSVD representation. It is also acceptable for $(I-R_m)P$ or $(I-P)R_m$ to consist of a low-rank part plus a localized part. Estimates for this spectral equivalence appear in Section A.3.

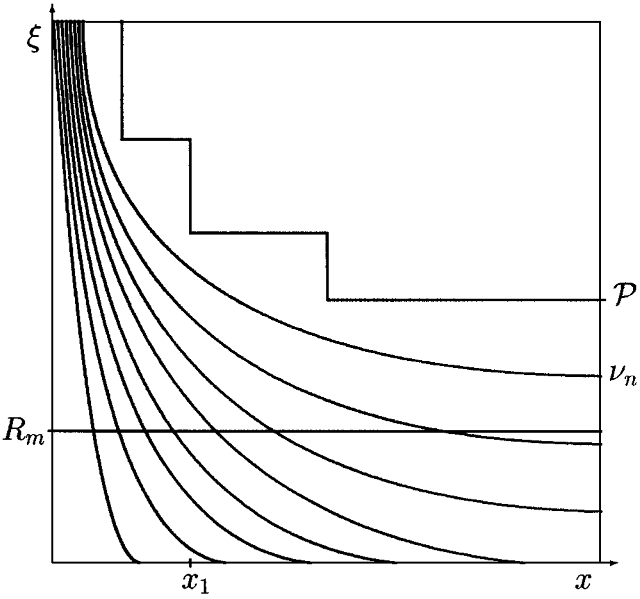

FIG. 4. Schematic of instantaneous frequency plots and the boundaries of the projectors $\mathcal{P}$ and R_m.

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 $-\nabla^2$ with periodic boundary conditions. Since the smallest eigenvalue of T is 0, $T_0 = T/\|T\| + \kappa^{-1} I$ 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 $\mathcal{O}(\log \kappa + \log\log(1/\epsilon))$ iterations, and the $\mathcal{O}(\log \kappa)$ 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 $-\nabla^2 - 300/|x|$ on the interval $[-1/2, 1/2]$ with periodic boundary conditions. In our first example, we choose $\mu = 0$ 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 approximation for the second derivative and arrange for the singularity in the potential to occur between sample points. The wavelets used were coiflets with seven vanishing moments.

EXAMPLE 4.1. First we use $\mu = 0$ ($n = 15$) and produce Fig. 5. Plotted is the number of coefficients used to represent the matrix as a fraction of the size of the full matrix, 512^2. The first plot is using standard coordinates and representing the matrix in sparse form, with truncation at 10^{-8}. Initially the matrix is banded and so sparse, but by the end of the iteration it is full. The second plot is again a sparse matrix representation with truncation at 10^{-8}, but is in wavelet coordinates. Initially it mimics the banded form of standard coordinates, but it remains sparse throughout the iteration, improving at the end. In the third plot we use the rough projector to reduce the size of the matrix and then use the sparse structure as before, in wavelet coordinates. The subspace we use for $\mathcal{P}$ is given by all wavelet coordinates that the eigenfunctions we wish to capture use at level 10^{-8}, and has dimension 322. In the true problem the eigenfunctions are not known, but this choice for $\mathcal{P}$ allows us to prove the principle without addressing the problem of the construction of $\mathcal{P}$ (see Section A.2). By using

TABLE I
Number of Iterations Required for the Cubic Sign Recursion (9) to Converge with 10^{-7} Accuracy

Condition number of T_0	Iterations
10^1	11
10^2	15
10^3	21
10^4	26

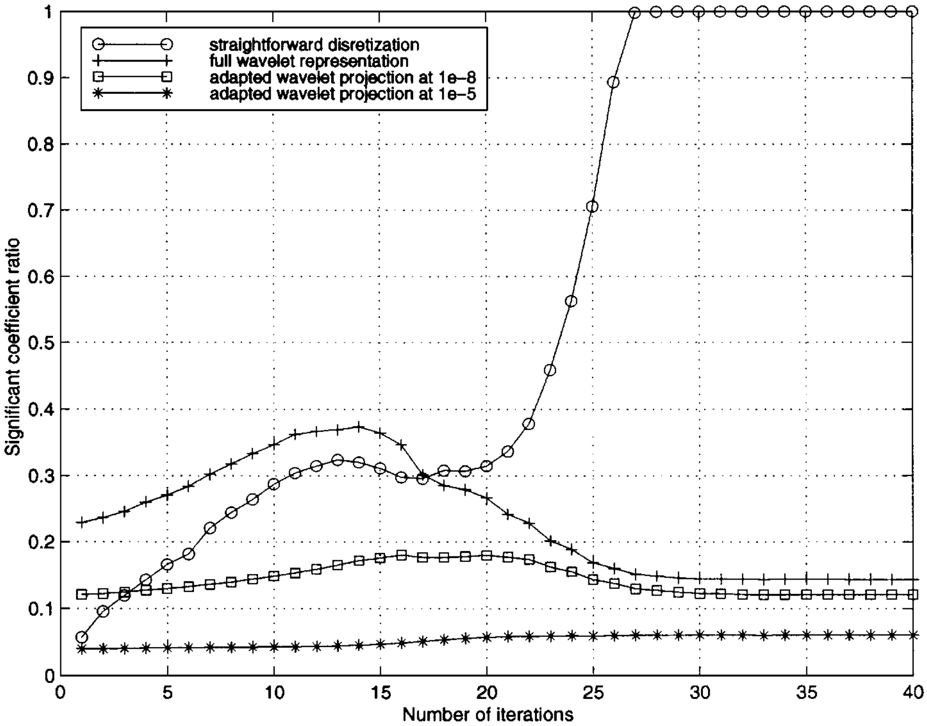

FIG. 5. (Example 4.1) Ratio of the number of significant coefficients above the threshold 10^{-8} to the total number of matrix elements during the cubic sign recursion (9).

this rough projector, we have eliminated the bump in the number of wavelet coefficients in the middle of the iteration. The fourth plot is similar to the third, except the rough projector is set to level 10^{-5}, giving $\mathcal{P}$ rank 133. The number of coefficients is cut further, without an increase in error, as discussed below. This final method gives about a factor of 16 savings in the size of the matrix and a factor of 64 in the number of computations.

In Fig. 6 we plot the error in the density for the methods in Fig. 5. This error is computed as the L^2 norm of the difference in the computed density from the actual density (computed from the eigenfunctions), divided by the L^2 norm of the actual density. We see convergence after 35 iterations, with some loss of accuracy due to conditioning ($\kappa \approx 5 \times 10^5$). The use of wavelet coordinates yields an extra digit of accuracy. The threshold used for the rough projector need not be tied to the condition number of the matrix. In this example, we are able to use a rough projector to only five digits without increasing the final error.

Example 4.2. In Fig. 7 we perform a similar test but with μ such that we capture 70 eigenfunctions. The first plot is using standard coordinates and truncating at 10^{-8}. The second plot uses wavelet coordinates with the same truncation. As predicted in Section 3.1, wavelet compression starts to fail, since the eigenfunctions are now more oscillatory. The rough projector at the 10^{-6} level is of full rank 512, and so provides no relief. The third plot is the number of coefficients needed for the PSVD representation at four levels of subdivision, including storage of the singular vectors, with truncation at 10^{-8}. We have not implemented the recursion in PSVD form, so we compute the PSVD from the full matrix constructed at higher precision at each iterate. (Truncation of the matrix tends to produce a large number of singular values just above the threshold.) As suggested in Section 3.3,

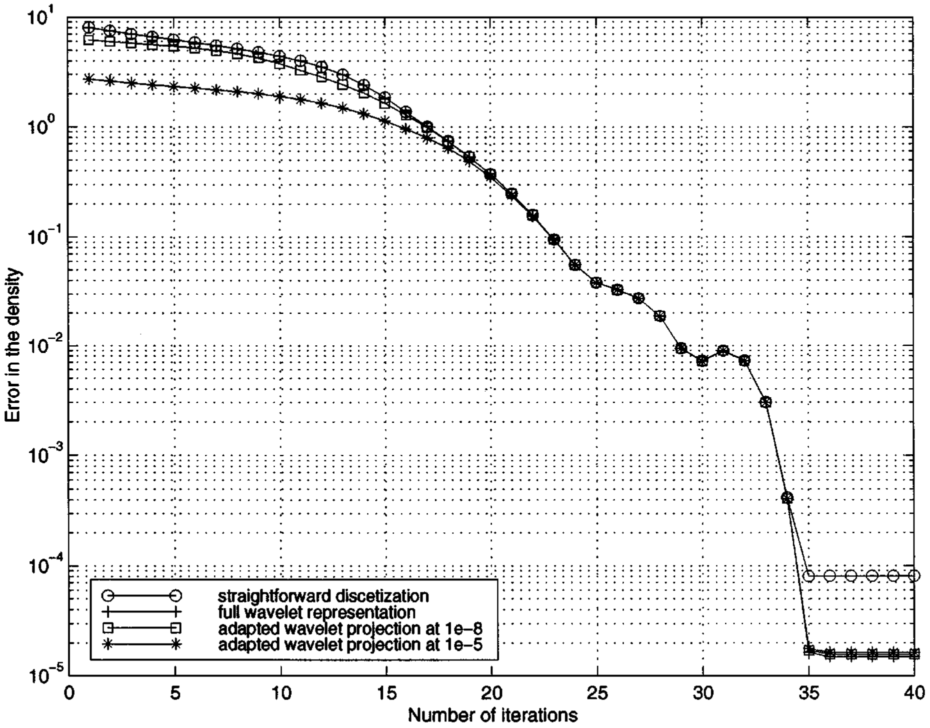

FIG. 6. (Example 4.1) Relative error in the density generated in Fig. 5.

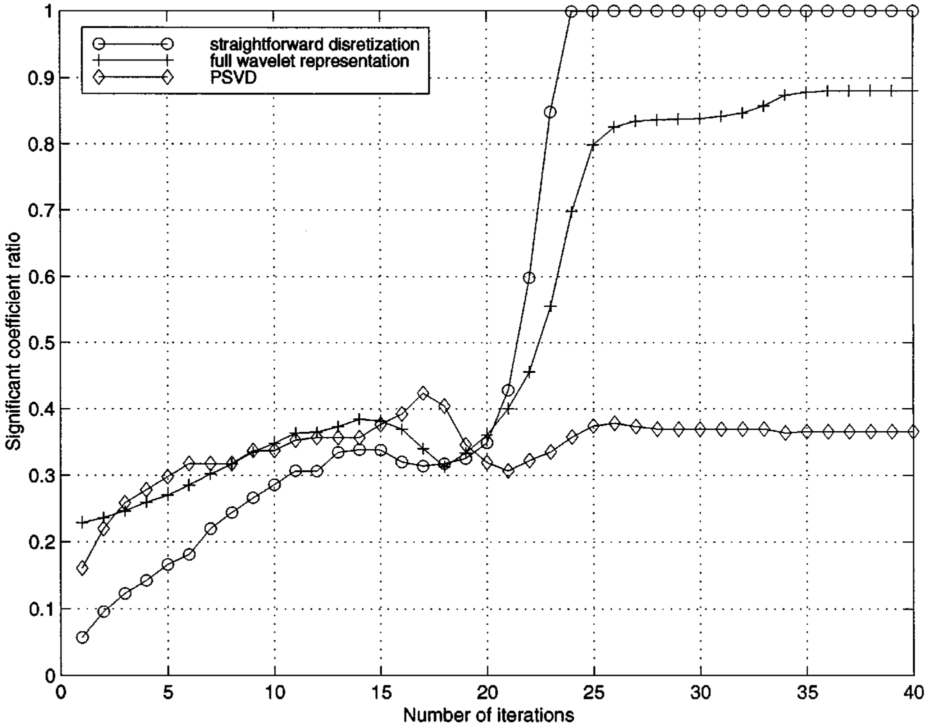

FIG. 7. (Example 4.2) Ratio of the number of significant coefficients above the threshold 10^{-8} to the total number of matrix elements during the recursion (9) for μ capturing 70 eigenvectors.

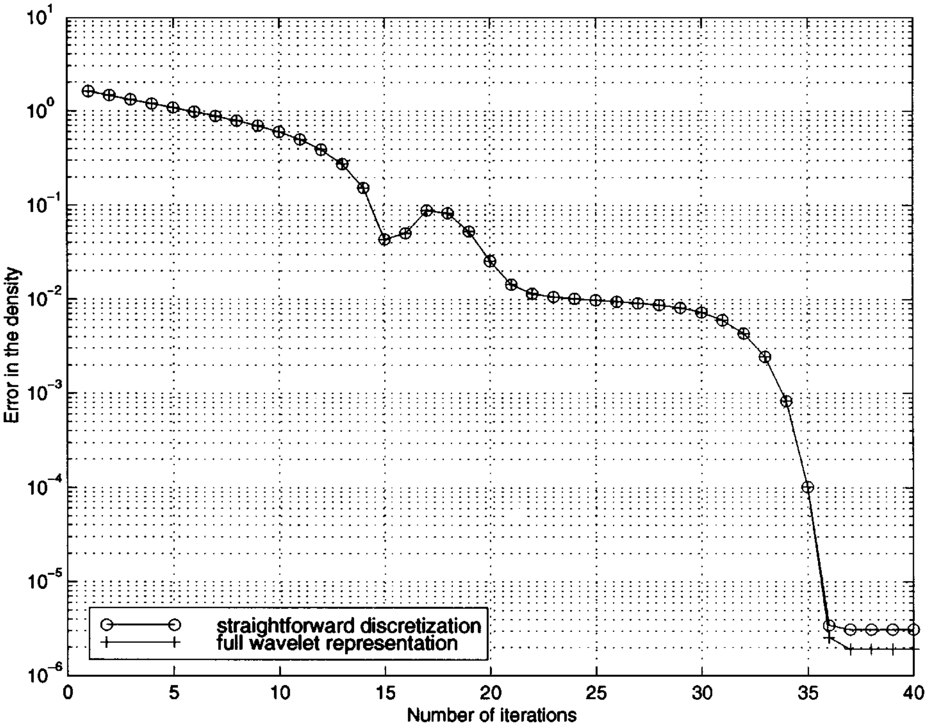

FIG. 8. (Example 4.2) Relative error in the density generated in Fig. 7.

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 $\kappa \approx 10^6$. We achieve slightly better performance than in Fig. 6 simply because we measure relative error.

5. EXTENSIONS

5.1. *Partitioning the Spectrum*

In this section we describe a method for partitioning the spectrum and indicate a situation when such a partition may be appropriate. We expect its greatest utility will be for other applications for spectral projectors, such as in wave propagation.

Let $\mu' < \mu < 0$. We will partition the spectrum into $(-\infty, \mu')$, (μ', μ), and (μ, ∞). Construct the spectral projector $P'(x, y)$ onto the interval $(-\infty, \mu')$ as above. Then construct

$$H' = (I - P')H(I - P'). \tag{19}$$

The eigenvectors corresponding to eigenvalues of H less than μ' now have zero eigenvalue with respect to H'. We can then project onto the (μ', μ) portion of the spectrum of H by computing

$$P'' = (I - \text{sign}(H' - \mu I))/2. \tag{20}$$

The electron density $\rho_n(x)$ is recovered by taking the diagonals of P'' and P' and adding.

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 $P'(x, y)$ and then $P''(x, y)$ 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' 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' 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 inertia for a symmetric matrix, but we demonstrate it with a simple example. This example also shows the effect of this projection on the eigenvectors. Suppose that $\mu = 0$, $\lambda_1 < \mu < \lambda_2$ and we have an operator

$$\begin{aligned} T &= \lambda_1 \psi_1(x)\psi_1(y) + \lambda_2 \psi_2(x)\psi_2(y) \\ &= \lambda_1(a_1\phi_1(x) + b_1\phi_2(x))(a_1\phi_1(y) + b_1\phi_2(y)) \\ &\quad + \lambda_2(a_2\phi_1(x) + b_2\phi_2(x))(a_2\phi_1(y) + b_2\phi_2(y)), \end{aligned} \tag{21}$$

where $\{\psi_1, \psi_2\}$ are the normalized eigenfunctions and $\{\phi_1, \phi_2\}$ is some basis. Assuming that $|b_1| < \epsilon$ we let $\mathcal{P}$ be the projection onto ϕ_1. Our operator becomes

$$\begin{aligned} \tilde{T} = \mathcal{P}T\mathcal{P} &= \lambda_1(a_1\phi_1(x))(a_1\phi_1(y)) + \lambda_2(a_2\phi_1(x))(a_2\phi_1(y)) \\ &= \left(\lambda_1 a_1^2\right)\phi_1(x)\phi_1(y) + \left(\lambda_2 a_2^2\right)\phi_1(x)\phi_1(y). \end{aligned} \tag{22}$$

In the subspace we wanted to keep, the eigenvalue λ_1 becomes $\lambda_1 a_1^2 = \lambda_1(1 - \mathcal{O}(\epsilon^2)) < 0$ and the eigenfunction $\psi_1(x)$ becomes $\phi_1(x) = (\psi_1(x) - \mathcal{O}(\epsilon)\phi_2(x))/(1 - \mathcal{O}(\epsilon^2))$. In the complementary subspace, λ_2 becomes $\lambda_2 a_2^2 \geq 0$.

A.2. *The Instantaneous Frequency Perspective*

In this section we examine how close $\mathcal{P}$ can be to P, justifying the arguments of Section 3.1. We use the local cosine basis and the quantitative quasi-classical method of [31] because they allow rigorous bounds. Local cosine is difficult to orthogonalize in higher dimensions, so we do not use it in practice.

A local cosine basis (see e.g., [18, 19]) is constructed as follows: We begin with a sequence of points on the line (interval, circle) $\cdots x_i < x_{i+1} \cdots$. Let $I_i = [x_i, x_{i+1}]$. We have a set of compatible bells, indexed by their interval, $\{b_i(x)\}$. These bells have the properties that $b_i(x)b_{i-1}(x)$ is even about x_i, $b_i(x)b_{i'}(x) = 0$ if $i' \neq i \pm 1$, and $\sum_i b_i^2(x) = 1$. On each interval we have a set of cosines of the proper scaling and shift, denoted

$$\left\{ c_i^p(x) = \sqrt{\frac{2}{x_{i+1} - x_i}} \cos\left(\frac{(p + 1/2)\pi(x - x_i)}{x_{i+1} - x_i}\right)\right\}_{p=0}^{\infty}. \tag{23}$$

The set $\{b_i(x)c_i^p(x)\}$ forms an orthonormal basis for the line.

The use of cosines instead of exponentials and the half-integer frequencies $p + 1/2$ is crucial for the orthogonality of these functions but has no effect on the decay estimates we need. To simplify the proof, we provide estimates using the exponentials

$$\left\{ e_i^p(x) = \sqrt{\frac{1}{x_{i+1} - x_i}} \exp\left(i \frac{p\pi(x - x_i)}{x_{i+1} - x_i}\right)\right\}_{p=0}^{\infty}. \tag{24}$$

For simplicity we also consider b_i to be supported on I_i, rather than overlapping with the neighboring interval.

In Section 3.1 we introduced the notion that those local cosine boxes that intersect the instantaneous frequency $\nu(x)$ should correspond to basis functions that yield significant coefficients (see Fig. 1). To make this notion rigorous, we can apply an integration by parts argument as in [31].

THEOREM A.1. *Let I be an interval of length l, ψ_i the normalized eigenfunction as above, $V^* \in \mathbf{R}$, and b_I a dilation and translation to I of a bell b on* $[0, 1]$; *then*

$$\left|\left\langle \psi_n, b_I e_I^p \right\rangle\right| \le \frac{l^2\|\lambda_n - V(x) - V^*\|_{L^\infty(I)} + \|b''\|_{L^\infty([0,1])} + 2|p\pi| \cdot \|b'\|_{L^\infty([0,1])}}{|l^2V^* - (p\pi)^2|}. \qquad (25)$$

We can manipulate (25) to determine for which p we have $|\langle \psi_n, b_I e_I^p\rangle| \le \epsilon$, where ϵ is the desired accuracy. If we choose the constant V^* to minimize $\|\lambda_n - V(x) - V^*\|_{L^\infty(I)}$, then $\sqrt{V^*}$ is the central frequency in the instantaneous frequency $\nu_n(x)$ for $\psi_n(x)$, and we can interpret our estimates on p as defining a band around $\nu_n(x)$ on which $\psi_n(x)$ lives. In Section A.3 we will use this band to show spectral equivalence. For the results of the current section, we are interested only in large p, so it would be sufficient to take $V^* = 0$.

Ignoring constant factors, we can conclude $|\langle \psi_n, b_I e_I^p\rangle| \le \epsilon$ when

$$p > \max\left\{ l\sqrt{V^* + \frac{\|\lambda_n - V(x) - V^*\|_{L^\infty(I)}}{\epsilon}},\ \sqrt{l^2V^* + \frac{\|b''\|_{L^\infty([0,1])}}{\epsilon}},\right.$$
$$\left. \frac{\|b'\|_{L^\infty([0,1])} + \sqrt{\|b'\|^2_{L^\infty([0,1])} + 2\epsilon^2 l^2 V^*}}{\epsilon} \right\}. \qquad (26)$$

The role of the second and third terms in (26) is to give a minimum value for p consistent with the uncertainty principle. Even in a region where $V(x)$ is nearly constant, the act of localizing with a bell implies that we will need a certain number of frequencies to represent ψ_n. We will ignore these two terms for our main analysis. In order to satisfy the first term in (26) it is sufficient to take

$$p > l\sqrt{V^*} + \frac{l\sqrt{\|\lambda_n - V(x) - V^*\|_{L^\infty(I)}}}{\epsilon}. \qquad (27)$$

Since $\sqrt{V^*}$ is the typical local frequency, the first term in (27) defines the basic sampling rate. The second term in (27) measures the "frequency spillage" due to the variation in $V(x)$. If λ_n is not available, we can use μ in (27). Applying this theorem on a properly chosen partition allows us to construct $\mathcal{P}$ as in Fig. 1.

We wish to calculate the total number of basis functions needed to construct $\mathcal{P}$ satisfying $\|\mathcal{P}P - P\| < \epsilon$. We will show that this number is the same order as the dimension of P, and thus $\mathcal{P}(I - P)$ does not have large rank. On any given partition, the contribution of the $l\sqrt{V^*}$ term in (27) to the number of basis functions needed, when summed over the intervals in the partition, takes the form of a Riemann sum which can be estimated by $\int_0^1 \sqrt{C/x}\, dx = 2\sqrt{C}$. On an interval $[x_i, x_{i+1}]$ of length l_i, the second term in (27) is bounded by $l_i\sqrt{l_i|V'(x_i)|}/\epsilon = (\sqrt{C}/\epsilon)l_i^{3/2}/x_i$. By choosing $l_i^{3/2} = \alpha x_i$ we make the number of basis functions needed per interval constant. Summing over the intervals, the total number can then be estimated by $\int_0^1 (\sqrt{C}/\epsilon)\alpha^{1/3}x^{-2/3}\, dx = (3\sqrt{C}/\epsilon)\alpha^{1/3}$. If we could choose $\alpha \to 0$ then this term would not contribute at all. The second and third terms in (26), however, bound α from below by a universal constant. These estimates really hold only on $(0, 1]$ since an interval which includes $x = 0$ would have $V^* = \infty$.

In order to improve the dependence on ϵ from ϵ^{-1} to $\log(1/\epsilon)$, one can apply the integration by parts technique used for Theorem A.1 k times and minimize over k. This technique fails only if $(|V^{(k)}(x)|)^{1/(k+1)}/k$ 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 $\mathcal{P}$ is representable in a local cosine subspace with dimension proportional to $\sqrt{C}$. The number of eigenfunctions n is also proportional to $\sqrt{C}$ (by, e.g., WKB estimates), so there is no fundamental obstruction to $\mathcal{P}$ closely matching P.

Proof of Theorem A.1. We will suppose our interval is $[0, l]$, so

$$\left\langle \psi_n, b_I e_I^p \right\rangle = \int_0^l \psi_n(x) b(x/l) \sqrt{\frac{1}{l}} \exp\left(i \frac{xp\pi}{l} \right) dx. \tag{28}$$

Integrating twice by parts, we obtain

$$= \int_0^l \left[(b''(x/l)/l^2)\psi_n(x) + 2(b'(x)/l)\psi_n'(x) + b(x)\psi_n''(x) \right] \sqrt{\frac{1}{l}} \frac{\exp(ixp\pi/l)}{-(p\pi/l)^2}\, dx. \tag{29}$$

Considering the $b'\psi'$ term separately and integrating by parts again, we have

$$= \frac{1}{-(p\pi/l)^2} \sqrt{\frac{1}{l}} \int_0^l \left[-\frac{b''(x)}{l^2} - (\lambda_n - V(x))b(x) + i2\frac{p\pi}{l}\frac{b'(x)}{l} \right] \exp\left(i\frac{xp\pi}{l} \right) \psi_n(x)\, dx. \tag{30}$$

Choosing any V^* we have

$$\begin{aligned} &\left\langle \psi_n, b_I e_I^p \right\rangle \left(1 + \frac{V^*}{-(p\pi/l)^2} \right) \\ &\quad = \frac{1}{-(p\pi/l)^2} \sqrt{\frac{1}{l}} \int_0^l \left[-\frac{b''(x)}{l^2} - (\lambda_n - V(x) - V^*)b(x) + i2\frac{p\pi}{l}\frac{b'(x)}{l} \right] \\ &\quad\quad \times \exp\left(i\frac{xp\pi}{l} \right) \psi_n(x)\, dx \end{aligned} \tag{31}$$

$$\begin{aligned} \left\langle \psi_n, b_I e_I^p \right\rangle = {} & \frac{1}{V^* - (p\pi/l)^2} \sqrt{\frac{1}{l}} \int_0^l \left[-\frac{b''(x)}{l^2} - (\lambda_n - V(x) - V^*)b(x) + i2\frac{p\pi}{l}\frac{b'(x)}{l} \right] \\ & \times \exp\left(i\frac{xp\pi}{l} \right) \psi_n(x)\, dx. \end{aligned} \tag{32}$$

Applying Hölder's inequality with the dual exponents $(1, \infty)$ and then the triangle inequality yields

$$\begin{aligned} \left| \left\langle \psi_n, b_I e_I^p \right\rangle \right| \le {} & \frac{1}{|V^* - (p\pi/l)^2|} \sqrt{\frac{2}{l}} \\ & \times \left(\left\| \frac{b''}{l^2} \right\|_\infty + \|\lambda_n - V(x) - V^*\|_\infty + 2\frac{p\pi}{l} \left\| \frac{b'}{l} \right\|_\infty \right) \|\psi_n(x)\|_{L^1([0,l])}. \end{aligned} \tag{33}$$

Since $\|\psi_n(x)\|_{L^2([0,1])} = 1$ we have $\|\psi_n(x)\|_{L^1([0,l])} \le \sqrt{l}$. Rearranging the l's yields the theorem. ■

A.3. *Spectral Equivalence*

For the techniques in Section 3.3 to be valid in our case, we need estimates to show that $-\nabla^2$ and $-\nabla^2 + V(x)$ 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 $|l^2 V^* - (p\pi)^2|^{-1}$. 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 $\{p : |\langle \psi_n, b_I e_I^p \rangle| > \epsilon\}$ than upper and lower bounds on this set, since the decay $V^* - (p\pi)^2$ is not symmetric about $\sqrt{V^*}$. Fixing $0 < x_1$, Theorem A.1 bounds the length of this set for the interval $[x_1, 1]$, $(l = 1 - x_1)$ by (ignoring constants)

$$\max\left\{\min\left\{l\sqrt{\frac{\|V(x) - V^*\|_\infty}{\epsilon}}, l\,\frac{\|V(x) - V^*\|_\infty}{\epsilon\sqrt{V^*}}\right\}, \frac{\|b'\|_\infty}{\epsilon}, \min\left\{\sqrt{\frac{\|b''\|_\infty}{\epsilon}}, \frac{\|b''\|_\infty}{\epsilon\sqrt{l^2 V^*}}\right\}\right\}. \tag{34}$$

As in (26) the second and third terms ensure that we do not violate the uncertainty principle when V is nearly constant. Since $p \approx lm$, we can translate the second part of the first term in (34) to a bound on $k - m$, where $R_k P \approx P$ and $P R_m \approx R_m$ on $[x_1, 1]$. This estimate then bounds the rank of $(I - R_m)R_k$ which in turn bounds the rank of $(I - R_m)P$. For the potential $-C/|x|$ we obtain

$$\frac{|C/x_1|}{\epsilon\sqrt{\mu - C/x_1}} = \frac{C}{\epsilon\sqrt{x_1}\sqrt{x_1\mu - C}}. \tag{35}$$

This bound goes to 0 as $\mu \to \infty$, but increases as $x_1 \to 0$. Using this method the other term $(I - P)R_m$ is negligible on $[x_1, 1]$.

We note finally that the Christoffel–Darboux summation formula is really a local result. We could partition $[0, 1]$ and apply bounds as above on each interval, obtaining a tighter result. Due to the difficulties of interpreting these results on the partition in Fig. 3, we will not pursue this idea. It does suggest, however, that the performance of this algorithm may be better than predicted via (35).

APPENDIX B: GREEN'S FUNCTION AND SPECTRAL PROJECTORS

One approach to constructing the electron density uses contour integrals,

$$\rho_n(x) = \frac{1}{2\pi i}\int_C G(x, x, z)\, dz, \tag{36}$$

where $G = (H - z)^{-1}$ and the contour C has the eigenvalues of interest in its interior. This approach shifts the difficulty from solving an eigenvalue problem at each step in the Kohn–Sham iteration to that of computing Green's function G and integrating over an appropriate contour to obtain $\rho_n(x)$. In this appendix we show that our sign function formulation is equivalent to the contour integral (36), namely

$$(I - \operatorname{sign}(T - \mu I))/2 = -\frac{1}{2\pi i}\int_C (T - zI)^{-1}\, dz, \tag{37}$$

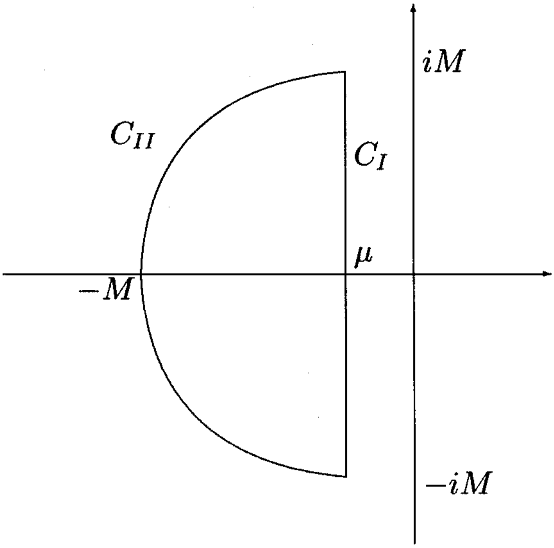

FIG. 9. The contour C, composed of C_{I} and C_{II}.

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

$$-\frac{1}{2\pi i}\int_C (T - zI)^{-1}\,dz = -\frac{1}{2\pi i}U^*\left(\int_C (D - zI)^{-1}\right) dzU, \tag{38}$$

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

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

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 through $-iM$ and connecting with the endpoints of C_{I}. We take the limit as $M \to \infty$.

First consider the integral over C_{II}. For M large enough, we may write the series representation

$$(\lambda - z)^{-1} = -z^{-1}\sum_{k=0}^{\infty}\left(\frac{\lambda}{z}\right)^k. \tag{40}$$

This series is uniformly convergent for all z on C_{II}. By writing $z = Me^{i\theta}$ and $dz = iMe^{i\theta}\,d\theta$, we obtain

$$-\frac{1}{2\pi i}\int_{C_{\mathrm{II}}} (\lambda - z)^{-1}\,dz = \frac{1}{2\pi}\sum_{k=0}^{\infty}\int_{-\pi/2}^{\pi/2}\left(\frac{\lambda}{Me^{i\theta}}\right)^k d\theta. \tag{41}$$

As $M \to \infty$ only the term $k = 0$ remains, leaving us with

$$\lim_{M\to\infty} -\frac{1}{2\pi i}\int_{C_{\mathrm{II}}} (\lambda - z)^{-1}\, dz = \frac{1}{2}. \tag{42}$$

For C_{I} we have $z = \mu + it$ and $dz = i\, dt$,

$$\lim_{M\to\infty} -\frac{1}{2\pi i}\int_{C_{\mathrm{I}}} (\lambda - z)^{-1}\, dz = \lim_{M\to\infty} -\frac{1}{2\pi}\int_{-M}^{M} \frac{\lambda - \mu + it}{(\lambda-\mu)^2 + t^2}\, dt. \tag{43}$$

The imaginary part of the integrand in (43) is odd, so the imaginary part of the integral is zero. For the real part, we have

$$\lim_{M\to\infty} -\frac{1}{2\pi}\int_{-M}^{M} \frac{\lambda - \mu}{(\lambda-\mu)^2 + t^2}\, dt = -\frac{1}{\pi}\lim_{M\to\infty} \tan^{-1}\left(\frac{M}{\lambda-\mu}\right) = -\frac{1}{2}\operatorname{sign}(\lambda - \mu), \tag{44}$$

and obtain (37).

REFERENCES

1. R. M. Dreizler and E. K. U. Gross, *Density Functional Theory: An Approach to the Quantum Many-Body Problem* (Springer-Verlag, Berlin/New York, 1990).
2. X-P. Li, R. W. Nunes, and D. Vanderbilt, Density-matrix electronic-structure method with linear system-size scaling, *Phys. Rev. B* **47**(16), 10891–10902 (1993).
3. W. Kohn, Density functional and density matrix method scaling linearly with the number of atoms, *Phys. Rev. Lett*. **76**(17), 3168–3171 (1996).
4. S. Goedecker and L. Colombo, Efficient linear scaling algorithms for tight-binding molecular dynamics, *Phys. Rev. Lett.* **73**(1), 122–125 (1994).
5. S. Goedecker and M. Teter, Tight-binding electronic-structure calculations and tight-binding molecular dynamics with localized orbitals, *Phys. Rev. B* **51**(15), 9455–9464 (1995).
6. R. Baer and M. Head-Gordon, Sparsity of the density matrix in Kohn–Sham density functional theory and an assessment of linear system-size scaling methods, *Phys. Rev. Lett.* (1997).
7. K. R. Bates, A. D. Daniels, and G. E. Scuseria, Comparison of conjugate gradient density matrix search and Chebychev expansion methods for avoiding diagonalization in large-scale electronic structure calculations, *J. Chem. Phys*. **109**(9), 3308–3312 (1998).
8. D. J. Kouri, Y. Huang, and D. K. Hoffman, Direct approach to density functional theory: Heaviside–Fermi level operator using a pseudopotential treatment, *J. Phys. Chem*. **100**, 7903–7910 (1996).
9. H. Röder, R. N. Silver, D. A. Drabold, and Jian Jun Dong, Kernel polynomial method for a nonorthogonal electronic-structure calculation of amorphous diamond, *Phys. Rev. B* **55**(23), 15382–15385 (1997).
10. C. S. Kenney and A. J. Laub, The matrix sign function, *IEEE Trans. Automat. Control* **40**(8), 1330–1348 (1995).
11. L. Auslander and A. Tsao, On parallelizable eigensolvers, *Adv. Appl. Math*. **13**, 253–261 (1992).
12. S. Huss-Lederman, A. Tsao, and T. Turnbull, *A Parallelizable Eigensolver for Real Diagonalizable Matrices with Real Eigenvalues*, Technical Report, PRISM Working Note No. 20.
13. P. Jones, J. Ma, and V. Rokhlin, A fast direct algorithm for the solution of the Laplace equation on regions with fractal boundaries, *J. Comput. Phys*. **113**(1) (1994).
14. T. Hrycak and V. Rokhlin, *An Improved Fast Multipole Alogrithm for Potential Fields*, Technical Report, YALEU/DCS/RR-1089, Yale University (1995).
15. V. Rokhlin and N. Yarvin, *A Generalized One-Dimensional Fast Multipole Method with Application to Filtering of Spherical Harmonics*, Technical Report, YALEU/DCS/RR-1142, Yale University (1998).

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).
25. R. A. Lippert, T. A. Arias, and A. Edelman, Multiscale computation with interpolating wavelets, *J. Comput. Phys*. **140**(2), 278–310 (1998).
26. G. Szegö, *Orthogonal Polynomials*, 4th ed. (Am. Math. Soc., Providence, RI, 1975).
27. B. Alpert, G. Beylkin, R. R. Coifman, and V. Rokhlin, Wavelet-like bases for the fast solution of second-kind integral equations, *SIAM J. Sci. Statist. Comput.* **14**(1), 159–174 (1993). [Technical Report, Department of Computer Science, Yale University, New Haven, CT, 1990]
28. B. Alpert, A class of bases in L^2 for the sparse representation of integral operators, *SIAM J. Math. Anal*. **24**(1), 246–262 (1993).
29. B. Alpert, G. Beylkin, D. Gines, and L. Vozovoi, *Adaptive Solution of Partial Differential Equations in Multiwavelet Bases* (1999), in progress.
30. G. Beylkin and R. Cramer, *A Multiresolution Approach to Fast Summation and Application of Singular Operators in Higher Dimensions* (1998), in progress.
31. M. J. Mohlenkamp, A fast transform for spherical harmonics, *J. Fourier Anal. Appl.* **5**(2) (1999).
32. R. Jakob–Chien and B. K. Alpert, A fast spherical filter with uniform resolution, *J. Comput. Phys.* **136**, 580–584 (1997).