

Discrete Symbol Calculus

Laurent Demanet
Department of Mathematics
Stanford University
450 Serra Mall
Stanford, CA 94305

Lexing Ying
Department of Mathematics
University of Texas at Austin
1 University Station/C1200
Austin, TX 78712

June 2008, revised June 2009

Abstract

This paper deals with efficient numerical representation and manipulation of differential and integral operators as symbols in phase-space, i.e., functions of space x and frequency ξ . The symbol smoothness conditions obeyed by many operators in connection to smooth linear partial differential equations allow to write fast-converging, non-asymptotic expansions in adequate systems of rational Chebyshev functions or hierarchical splines. The classical results of closedness of such symbol classes under multiplication, inversion and taking the square root translate into practical iterative algorithms for realizing these operations directly in the proposed expansions. Because symbol-based numerical methods handle operators and not functions, their complexity depends on the desired resolution N very weakly, typically only through $\log N$ factors. We present three applications to computational problems related to wave propagation: 1) preconditioning the Helmholtz equation, 2) decomposing wavefields into one-way components and 3) depth extrapolation in reflection seismology. The software can be downloaded from the software section of either author's website.

Acknowledgements. The first author is partially supported by an NSF grant. The second author is partially supported by an NSF grant, a Sloan Research Fellowship, and a startup grant from the University of Texas at Austin.

1 Introduction

There are many interesting puzzles related to algorithmic complexity and scalability left in relation to partial differential equations. Some of these questions are idealized versions of challenges encountered in industrial applications. One notable success story where mathematics played a role is the fast multipole method of Greengard and Rokhlin [22], now an authoritative algorithmic tool in electrodynamics. There, the problem was to provide a $O(N)$ algorithm for computing the electrostatic interaction between all pairs among N charged particles.

Of growing interest is the related question:

*If required to solve the same linear problem thousands of times,
can an adequate precomputation lower the overall algorithmic complexity?*

For instance, in the scope of the boundary integral formulations of electromagnetism, addressing this question would mean going beyond fast multipole or other methods for the integral kernels,

and instead precomputing the whole linear map that solves the integral equation (for the source density in terms of the incident fields.) A satisfactory answer would be a fast applicator for this linear map, with algorithmic complexity independent of the frequency of the incoming fields. The authors are unaware that any progress has been made on this question yet.

This paper deals with another instance of “computational preparation” for solving a problem multiple times, this time in the scope of simple linear partial differential equations in variable, smooth media. Throughout this paper, we will be interested in the efficient representation of *functions* of elliptic operators, such as

$$A = I - \operatorname{div}(\alpha(x)\nabla),$$

where $\alpha(x) > c > 0$ is smooth, and for simplicity, $x \in [0, 1]^d$ with periodic boundary conditions. The inverse, the square root, and the exponential of operators akin to A are operations that all play important roles in the inverse problem of reflection seismology. There are of course other application areas than elastic wave propagation, but seismology is a problem of wide interest, where the same equations have to be solved literally thousands of times, and where the physics lies in the spatial variability of the elastic parameters — not in the boundary conditions.

Most numerical methods for inverting A , say, would manipulate a right-hand side until convergence, and leverage sparsity of a matrix realization of A in doing so. Large condition numbers for A may considerably slow down convergence. Forming A^{-1} directly would be a way to avoid these iterations, but discretizations of integral operators rarely give rise to sparse matrices. Instead, we present expansion schemes and iterative algorithms for manipulating functions of A as “symbols”, whose numerical realization make little or no reference to the functions of x to which these operators may later be applied.

The central question is that of choosing a representation that will be computationally advantageous over wide classes of differential and integral operators. If a function $f(x)$ has N degrees of freedom — if for instance it is sampled on N points — then a direct representation of operators acting on $f(x)$ would in general require N^2 degrees of freedom. There are many known methods for bringing down this count to $O(N)$ or $O(N \log N)$ in specific cases, such as leveraging sparsity, computing convolutions via FFT, low-rank approximations, fast summation methods [22, 25], wavelet or x-let expansions [3], partitioned SVD and H-matrices [9, 24], semi-separable matrices [23], butterfly algorithms [38], etc.

The framework presented in this paper is different: the algorithmic complexity of representing a discrete symbol is at most logarithmic in N , at least for operators belonging to certain standard classes.

In the spirit of work by Beylkin [4], Hackbusch [24], and others, the symbol representation is then used for perform numerical *operator calculus*. Composition of two operators is the main operation that requires a low-level implementation, aware of how symbols are realized. Once composition is available, it is the building block for iterations that compute functions of A without ever forming products Af — hence the name operator calculus. The symbol representation provides the required degree of compression for performing this calculus efficiently, and almost entirely circumvents complexity overheads associated to ill-conditioning. The complexity for all calculus operations is at most $O(\log^2 N)$, with a constant that depends logarithmically on the condition number. It is only when applying an operator to a function that the complexity is superlinear in N , in our case $O(N \log N)$, independently of the condition number.

A large fraction of the paper is devoted to covering two different symbol expansion schemes, as well as the numerical realization of the calculus operations. The computational experiments involving the inverse validate the potential of discrete symbol calculus over a standard preconditioner for

solving a simple test elliptic problem. Other computational experiments involve the square root: it is shown that this operation can be done accurately without forming the full matrix.

The main practical limitations of discrete symbols calculus currently are 1) the inability to properly handle boundary conditions, and 2) the quick deterioration of performance in nonsmooth media.

1.1 Smooth symbols

Let A denote a generic differential or singular integral operator acting on functions of $x \in \mathbb{R}^d$, with kernel representation

$$Af(x) = \int k(x, y) f(y) dy, \quad x, y \in \mathbb{R}^d.$$

Expanding the distributional kernel $k(x, y)$ in some basis would be cumbersome because of the presence of a singularity along the diagonal $x = y$. For this reason we choose to consider operators as *pseudodifferential symbols* $a(x, \xi)$, by considering their action on the Fourier transform¹ $\hat{f}(\xi)$ of $f(x)$;

$$Af(x) = \int e^{2\pi i x \cdot \xi} a(x, \xi) \hat{f}(\xi) d\xi.$$

We typically reserve uppercase letters for operators, except for occasionally using the analysts' attractive notation

$$A = a(x, D), \quad \text{where} \quad D = -\frac{i}{2\pi} \nabla_x.$$

By passing to the ξ variable, the singularity of $k(x, y)$ along $x = y$ is turned into the oscillating factor $e^{2\pi i x \cdot \xi}$ regardless of the type of singularity. This factor contains no information and is naturally discounted by focusing on the non-oscillatory symbol $a(x, \xi)$.

Symbols $a(x, \xi)$ are not merely C^∞ functions of x and ξ , but their smoothness properties are nevertheless well understood by mathematicians [46]. A symbol defined on $\mathbb{R}^d \times \mathbb{R}^d$ is said to be pseudodifferential of order m and type (ρ, δ) if it obeys

$$|\partial_\xi^\alpha \partial_x^\beta a(x, \xi)| \leq C_{\alpha\beta} \langle \xi \rangle^{m-\rho|\alpha|+\delta|\beta|}, \quad \text{where} \quad \langle \xi \rangle \equiv (1 + |\xi|^2)^{1/2}, \quad (1)$$

for all multi-indices α, β . In this paper, we mostly consider the special case of the type $(1, 0)$:

$$|\partial_\xi^\alpha \partial_x^\beta a(x, \xi)| \leq C_{\alpha\beta} \langle \xi \rangle^{m-|\alpha|}, \quad (2)$$

which is denoted S^m . An operator whose symbol $a \in S^m$ belongs by definition to the class Ψ^m .

The main feature of symbols in S^m is that *the larger* $|\xi|$, *the smoother the symbol in* ξ . Indeed, one power of $\langle \xi \rangle$ is gained for each differentiation. For instance, the symbols of differential operators of order m are polynomials in ξ and obey (2) when they have C^∞ coefficients. Large classes of singular integral operators also have symbols in the class S^m [46].

The standard treatment of pseudodifferential operators makes the further assumption that some symbols can be represented as polyhomogeneous series, such as

$$a(x, \xi) \sim \sum_{j \geq 0} a_j(x, \arg \xi) |\xi|^{m-j}, \quad (3)$$

¹Our conventions in this paper:

$$\hat{f}(\xi) = \int e^{-2\pi i x \cdot \xi} f(x) dx. \quad f(x) = \int e^{2\pi i x \cdot \xi} \hat{f}(\xi) d\xi.$$

which defines the “classical” symbol class S_{cl}^m when the a_j are of class C^∞ . The corresponding operators are said to be in the class Ψ_{cl}^m . The series should be understood as an asymptotic expansion; it converges only when adequate cutoffs smoothly removing the origin multiply each term. Only then, the series does not converge to $a(x, \xi)$, but to an approximation that differs from a by a smoothing remainder $r(x, \xi)$, smoothing in the sense that $|\partial_\xi^\alpha \partial_x^\beta r(x, \xi)| = O(\langle \xi \rangle^{-\infty})$. For instance, an operator is typically transposed, inverted, etc. modulo a smoothing remainder [29].

The subclass (3) is central for applications to partial differential equations (PDE) — it is the cornerstone of theories such as geometrical optics — but the presence of remainders is a nonessential feature that should be avoided in the design of efficient numerical methods. The lack of convergence in (3) may be acceptable in the course of a mathematical argument, but it takes great additional effort to turn such series into accurate numerical methods; see [47] for an example. The objective of this paper is to find adequate substitutes for (3) that promote asymptotic series into fast-converging expansions.

It is the behavior of symbols at the origin $\xi = 0$ that makes S^m and S_{cl}^m only adequate for large- $|\xi|$ asymptotic analysis. In practice, exact symbols may have a singular behavior at $\xi = 0$. This issue should not be a distraction: the important feature is that the symbol should be smooth in ξ far away from the origin, and it is very robust. We will have more to say on the proper numerical treatment of $\xi = 0$ in the sequel.

There are in general no explicit formulas for the symbol of a function of an operator. Fortunately, some results in the literature guarantee exact closedness of the symbol classes (2) or (3) under inversion and taking the square root, without smoothing remainders. A symbol $a \in S^m$, or an operator $a(x, D) \in \Psi^m$, is said to be elliptic when there exists $R > 0$ such that

$$|a^{-1}(x, \xi)| \leq C |\xi|^{-m}, \quad \text{when} \quad |\xi| \geq R.$$

- It is a basic result that if $A \in \Psi^{m_1}, B \in \Psi^{m_2}$, then $AB \in \Psi^{m_1+m_2}$. See for instance Theorem 18.1.8 in [29], Volume 3.
- It is also a standard fact that if $A \in \Psi^m$, then its adjoint $A^* \in \Psi^m$.
- If $A \in \Psi^m$, and A is elliptic and invertible² on L^2 , then $A^{-1} \in \Psi^{-m}$. This result was proved by Shubin in 1978 in [44].
- For the square root, we also assume ellipticity and invertibility. It is furthermore convenient to consider operators on compact manifolds, in a natural way through Fourier transforms in each coordinate patch, so that they have discrete spectral expansions. A square root $A^{1/2}$ of an elliptic operator A with spectral expansion $A = \sum_j \lambda_j E_j$, where E_j are the spectral projectors, is simply

$$A^{1/2} = \sum_j \lambda_j^{1/2} E_j, \tag{4}$$

with of course $(A^{1/2})^2 = A$. In 1967, Seeley [42] studied such expressions for elliptic $A \in \Psi_{\text{cl}}^m$, in the context of a much more general study of complex powers of elliptic operators. If in addition m is an even integer, and an adequate choice of branch cut is made in the complex plane, then Seeley showed that $A^{1/2} \in \Psi_{\text{cl}}^{m/2}$; see [45] for an accessible proof that involves the complex contour “Dunford” integral reformulation of (4).

²In the sense that A is a bijection from $H^m(\mathbb{R}^d)$ to $L^2(\mathbb{R}^d)$, hence obeys $\|Af\|_{L^2} \leq C\|f\|_{H^m}$. Ellipticity, in the sense in which it is defined for symbols, obviously does not imply invertibility.

We do not know of a corresponding closedness result under taking the square root, for the non-classical class Ψ^m . In practice, one is interested in manipulating operators that come from PDE on bounded domains with certain boundary conditions; the extension of the theory of pseudodifferential operators to bounded domains is a difficult subject that this paper has no ambition of addressing. Let us also mention in passing that the exponential of an elliptic, non-self-adjoint pseudodifferential operator is not in general itself pseudodifferential.

Numerically, it is easy to check that smoothness of symbols is remarkably robust under inversion and taking the square root of the corresponding operators, as the following simple one-dimensional example shows.

Let $A := 4\pi^2 I - \operatorname{div}(\alpha(x)\nabla)$ where $\alpha(x)$ is a random periodic function over the periodized segment $[0, 1]$, essentially bandlimited as shown in Figure 1(a). The symbol of this operator is

$$a(x, \xi) = 4\pi^2(1 + \alpha(x)|\xi|^2) - 2\pi i \nabla \alpha(x) \cdot \xi,$$

which is of order 2. In Figure 1(b), we plot the values of $a(x, \xi)\langle \xi \rangle^{-2}$ for x and ξ on a Cartesian grid.

Since A is elliptic and invertible, its inverse $C = A^{-1}$ and square root $D = A^{1/2}$ are both well defined. Let use $c(x, \xi)$ and $d(x, \xi)$ to denote their symbols. From the above theorems, we know that the orders of $c(x, \xi)$ and $d(x, \xi)$ are respectively -2 and 1 . We do not believe that explicit formulae exist for these symbols, but the numerical values of $c(x, \xi)\langle \xi \rangle^2$ and $d(x, \xi)\langle \xi \rangle^{-1}$ are shown in Figure 1(c) and (d), respectively. These plots demonstrate regularity of these symbols in x and in ξ ; observe in particular the disproportionate smoothness in ξ for large $|\xi|$, as predicted by the class estimate (2).

1.2 Symbol expansions

Figure 1 suggests that symbols are not only smooth, but that they should be highly *separable* in x vs. ξ . We therefore use expansions of the form

$$a(x, \xi) = \sum_{\lambda, \mu} a_{\lambda, \mu} e_{\lambda}(x) g_{\mu}(\xi) \langle \xi \rangle^{d_a}, \quad (5)$$

where e_{λ} and g_{μ} are to be determined, and $\langle \xi \rangle^{d_a} \equiv (1 + |\xi|^2)^{d_a/2}$ encodes the order d_a of $a(x, \xi)$. This choice is in line with recent observations of Beylkin and Mohlenkamp [5] that functions and kernels in high dimensions should be represented in separated form. In this paper we have chosen to focus on two-dimensional x , i.e. $(x, \xi) \in \mathbb{R}^4$, which is already considered high-dimensional by the standard of numerical analysts. The curse of dimensionality would make unpractical any fine Cartesian sampling in four dimensions.

The functions $e_{\lambda}(x)$ and $g_{\mu}(\xi)$ should be chosen such that the interaction matrix $a_{\lambda, \mu}$ is as small as possible after accurate truncation. Their choice also depends on the domain over which the operator is considered. In what follows we will assume that the x -domain is the periodized unit square $[0, 1]^2$ in two dimensions. Accordingly, it makes sense to take for $e_{\lambda}(x)$ the complex exponentials $e^{2\pi i x \cdot \lambda}$ of a Fourier series. The choice of $g_{\mu}(\xi)$ is more delicate, as x and ξ do not play symmetric roles in the estimate (2). In short, we need adequate basis functions for smooth functions on \mathbb{R}^2 that behave like a polynomial of $1/|\xi|$ as $\xi \rightarrow \infty$, and otherwise present smooth angular variations. We present two solutions:

- A *rational Chebyshev interpolant*, where $g_{\mu}(\xi)$ are complex exponentials in angle $\theta = \arg \xi$, and scaled Chebyshev functions in $|\xi|$, where the scaling is an algebraic map $s = \frac{|\xi| - L}{|\xi| + L}$. More details are in Section 2.1.




Figure 1: Smoothness of the symbol in ξ . (a) The coefficient $\alpha(x)$. (b) $a(x, \xi) \langle \xi \rangle^{-2}$ where $a(x, \xi)$ is the symbol of A . (c) $c(x, \xi) \langle \xi \rangle^2$ where $c(x, \xi)$ is the symbol of $C = A^{-1}$. (d) $d(x, \xi) \langle \xi \rangle^{-1}$ where $d(x, \xi)$ is the symbol of $D = A^{1/2}$.

- A *hierarchical spline interpolant*, where $g_\mu(\xi)$ are spline functions with control points placed in a multiscale way in the frequency plane, in such a way that they become geometrically scarcer as $|\xi| \rightarrow \infty$. More details are in Section 2.2.

Since we are considering x in the periodized square $[0, 1]^2$, the Fourier variable ξ is restricted to having integer values, i.e., $\xi \in \mathbb{Z}^2$, and the Fourier transform should be replaced by a Fourier series. Pseudodifferential operators are then defined through

$$a(x, D)f(x) = \sum_{\xi \in \mathbb{Z}^2} e^{2\pi i x \cdot \xi} a(x, \xi) \hat{f}(\xi), \quad (6)$$

where $\hat{f}(\xi)$ are the Fourier series coefficients of f . There is nothing essential about ξ being discrete in this formula: it is still the smoothness of the underlying functions of $\xi \in \mathbb{R}^2$ that dictates the convergence rate of the proposed expansions.

The following results quantify the performance of the two approximants introduced above. We refer to an approximant \tilde{a} as being truncated to M terms when all but at most M elements are put to zero in the interaction matrix $a_{\lambda, \mu}$ in (5).

Theorem 1. (Rational Chebyshev approximants). *Assume that $a \in S_{cl}^m$, that it is properly supported, and assume furthermore that the a_j in equation (3) have tempered growth, in the sense*

that there exists $Q, R > 0$ such that

$$|\partial_\theta^\alpha \partial_x^\beta a_j(x, \theta)| \leq Q_{\alpha, \beta} \cdot R^j. \quad (7)$$

Denote by \tilde{a} the rational Chebyshev expansion of a (introduced in Section 2.1), properly truncated to M terms. Call \tilde{A} and A the corresponding pseudodifferential operators on $L^2([0, 1]^2)$, defined by (6). Then, there exists a choice of M obeying

$$M \leq C_n \cdot \varepsilon^{-1/n}, \quad \forall n > 0,$$

such that

$$\|\tilde{A} - A\|_{H^m([0, 1]^2) \rightarrow L^2([0, 1]^2)} \leq \epsilon.$$

Theorem 2. (Hierarchical spline approximants). Assume that $a \in S^m$, and that it is properly supported. Denote by \tilde{a} the expansion of a in hierarchical splines for ξ (introduced in Section 2.2), and in a Fourier series for x , properly truncated to M terms. Call \tilde{A} and A the corresponding pseudodifferential operators on $L^2([0, 1]^2)$, defined by (6). Introduce P_N the orthogonal projector onto frequencies obeying

$$\max(|\xi_1|, |\xi_2|) \leq N.$$

Then there exists a choice of M obeying

$$M \leq C \cdot \varepsilon^{-2/(p+1)} \cdot \log N,$$

where p is the order of the spline interpolant, and for some $C > 0$, such that

$$\|(\tilde{A} - A)P_N\|_{H^m([0, 1]^2) \rightarrow L^2([0, 1]^2)} \leq \epsilon.$$

The important point of these theorems is that M is either constant in N (Theorem 1), or grows like $\log N$ (Theorem 2), where N is the bandlimit of the functions to which the operator is applied.

1.3 Symbol operations

At the level of kernels, composition of operators is a simple matrix-matrix multiplication. This property is lost when considering symbols, but composition remains simple enough that the gains in dealing with small interaction matrices $a_{\lambda, \mu}$ as in (5) are far from being offset.

The *twisted product* of two symbols a and b , is the symbol of their composition. It is defined as $(a \# b)(x, D) = a(x, D)b(x, D)$ and obeys

$$a \# b(x, \xi) = \int \int e^{-2\pi i(x-y) \cdot (\xi - \eta)} a(x, \eta) b(y, \xi) dy d\eta.$$

This formula holds for $\xi, \eta \in \mathbb{R}^d$, but in the case when frequency space is discrete, the integral in η is to be replaced by a sum. In Section 3 we explain how to evaluate this formula efficiently using the symbol expansions discussed earlier.

Textbooks on pseudodifferential calculus describe asymptotic expansions of $a \# b$ where negative powers of $|\xi|$ are matched at infinity [29, 20, 45]. As alluded to previously, we are not interested in making simplifications of this kind.

Composition can be regarded as a building block for performing many other operations using iterative methods. Functions of operators can be computed by substituting the twisted product for the matrix-matrix product in any algorithm that computes the corresponding function of a matrix. For instance,

- The inverse of a positive-definite operator can be obtained via a Neumann iteration, or via a Schulz iteration;
- There exist many choices of iterations for computing the square root and the inverse square root of a matrix [28], such as the Schulz-Higham iteration;
- The exponential of a matrix can be obtained by the scaling-and-squaring method; etc.

These examples are discussed in detail in Section 3.

Two other operations that resemble composition from the algorithmic viewpoint, are 1) transposition, and 2) the Moyal transform for passing to the Weyl symbol. They are also discussed below.

Lastly, this work would be incomplete without a routine for applying a pseudodifferential operator to a function, from the knowledge of its symbol. The type of separated expansion considered in equation (5) suggests a very simple algorithm for this task³, detailed in Section 3.

1.4 Applications

It is natural to apply discrete symbol calculus to the numerical solutions of linear PDE with variable coefficients. We outline several examples in this section, and report on the numerical results in Section 4.

In all of these applications, the solution takes two steps. First, discrete symbol calculus is used to construct the symbol of the operator that solves the PDE problem. Since “data” like a right hand side, initial conditions, or boundary conditions have not been queried yet, the computational cost of this step is mostly independent of the size of the data. Once the operator is ready in its symbol form, we apply the operator to the data in the second step.

The two regimes in which this approach could be preferred is when either 1) the complexity of the medium (coefficient in the PDE) is low compared to the complexity of the data, or 2) the PDE needs to be solved so many times that a precomputation step becomes beneficial.

A first, toy application of discrete symbol calculus is to the numerical solution of the simplest elliptic PDE,

$$Au := (I - \operatorname{div}(\alpha(x)\nabla)u = f \tag{8}$$

with $\alpha(x) > 0$, and periodic boundary conditions on a square. If $\alpha(x)$ is a constant function, the solution requires only two Fourier transforms, since the operator is diagonalized by the Fourier basis. For variable $\alpha(x)$, discrete symbol calculus can be seen as a natural generalization of this fragile Fourier diagonalization property: we construct the symbol of A^{-1} directly, and once the symbol of A^{-1} is ready, applying it to the function f requires only a small number of Fourier transforms.

The second application of discrete symbol calculus concerns the Helmholtz equation

$$Lu := \left(-\Delta - \frac{\omega^2}{c^2(x)} \right) u = f(x) \tag{9}$$

where the sound speed $c(x)$ is a smooth function in x , in a periodized square. The numerical solution of this problem is quite difficult since the operator L is not positive definite. Efficient techniques such as multigrid cannot be used directly for this problem; a discussion can be found in [15]. A

³This part is not original; it was already considered in previous work by Emmanuel Candès and the authors in [11], where the more general case of Fourier integral operators was considered. See also [1].

standard iterative algorithm, such as MINRES or BIGGSTAB, can easily take tens of thousands of iterations to converge. One way to obtain faster convergence is to solve a preconditioned system

$$M^{-1}Lu = M^{-1}f \quad (10)$$

with

$$M := -\Delta + \frac{\omega^2}{c^2(x)} \quad \text{or} \quad M := -\Delta + (1 + i)\frac{\omega^2}{c^2(x)}$$

Now at each iteration of the preconditioned system, we need to invert a linear system for the preconditioner M . Multigrid is typically used for this step [16], but discrete symbol calculus offers a way to directly precompute the symbol of M^{-1} . Once it is ready, applying M^{-1} to a function at each iteration is reduced to a small number of Fourier transforms—three or four when $c(x)$ is very smooth—which we anticipate to be very competitive versus a multigrid method.

Another important application of the discrete symbol calculus is to *polarizing* the initial condition of a linear hyperbolic system. Let us consider the following variable coefficient wave equation on the periodic domain $x \in [0, 1]^2$,

$$\begin{cases} u_{tt} - \operatorname{div}(\alpha(x)\nabla u) = 0 \\ u(0, x) = u_0(x) \\ u_t(0, x) = u_1(x) \end{cases} \quad (11)$$

with the extra condition $\int u_1(x)dx = 0$. The operator $L := -\operatorname{div}(\alpha(x)\nabla)$ is symmetric positive definite; let us define P to be its square root $L^{1/2}$. We can then use P to factorize the wave equation as

$$(\partial_t + iP)(\partial_t - iP)u = 0.$$

As a result, the solution $u(t, x)$ can be represented as

$$u(t, x) = e^{itP}u_+(x) + e^{-itP}u_-(x)$$

where the polarized components $u_+(x)$ and $u_-(x)$ of the initial condition are given by

$$u_+ = \frac{u_0 + (iP)^{-1}u_1}{2} \quad \text{and} \quad u_- = \frac{u_0 - (iP)^{-1}u_1}{2}.$$

To compute u_+ and u_- , we first use discrete symbol calculus to construct the symbol of P^{-1} . Once the symbol of P^{-1} is ready, the computation of u_+ and u_- requires only applying P^{-1} to the initial condition. Computing e^{itP} is a difficult problem that we do not address in this paper.

Finally, discrete symbol calculus has a natural application to the problem of depth extrapolation, or migration, of seismic data. In the Helmholtz equation

$$\Delta_{\perp} + \frac{\partial^2 u}{\partial z^2} + \frac{\omega^2}{c^2(x, z)}u = F(x, z, k),$$

we can separate the Laplacian as $\Delta = \Delta_{\perp} + \frac{\partial^2}{\partial z^2}$, and factor the equation as

$$\left(\frac{\partial}{\partial z} - B(z) \right) v = F(x, z, k) - \frac{\partial B}{\partial z}(z)u, \quad \left(\frac{\partial}{\partial z} + B(z) \right) u = v \quad (12)$$

where $B = \sqrt{-\Delta_{\perp} - \omega^2/c^2(x, z)}$ is called the one-way wave propagator, or single square root (SSR) propagator. We may then focus on the equation for v , called the SSR equation, and solve it for

decreasing z from $z = 0$. The term $\frac{\partial B}{\partial z}(z)u$ above is sometimes neglected, as we do in the sequel, on the basis that it introduces no new singularities.

The symbol of B^2 is not elliptic; its zero level set presents a well-known issue with this type of formulation. In Section 4, we introduce an adequate “directional” cutoff strategy for removing the singularities that would otherwise appear, hence neglecting turning rays and evanescent waves. Discrete symbol calculus is then used to compute a proper operator square root. We show how to solve the SSR equation approximately using an operator exponential of B , also realized via discrete symbol calculus. Unlike traditional methods of seismic imaging (discussed in Section 1.6 below), the *only* simplification we make here is the directional cutoff just mentioned.

1.5 Harmonic analysis of symbols

It is instructive to compare the symbol expansions of this paper with another type of expansion thought to be efficient for smooth differential and integral operators, namely wavelets.

Consider $x \in [0, 1]$ for simplicity. The standard matrix of an operator A in a basis of wavelets $\psi_{j,k}(x) = 2^{j/2}\psi(2^j x - n)$ of $L^2([0, 1])$ is simply $\langle \psi_{j,k}, A\psi_{j',k'} \rangle$. Such wavelet matrices were first considered by Meyer in [37], and later by Beylkin, Coifman, and Rokhlin in [3], for the purpose of obtaining sparse expansions of singular integral operators in the Calderón-Zygmund class. Their result is that either $O(N)$ or $O(N \log N)$ elements suffice to represent a N -by- N matrix accurately, in the ℓ_2 sense, in a wavelet basis. This result is not necessarily true in other bases such as Fourier series or local cosines, and became the starting point of much activity in some numerical analysis circles in the 1990s.

In contrast, the expansions proposed in this paper assume a class of operators with symbols in the S^m class defined in (2), but achieve accurate compression with $O(1)$ or $O(\log N)$ elements. This stark difference is illustrated in Figure 2.

With symbols, tasks such as inversion and computing the square root are realized in $O(\log^2 N)$ operations, still way sublinear in N . It is only when the operator needs to be applied to functions defined on N points, as a “post-computation”, that the complexity becomes $C \cdot N \log N$. This constant C is proportional to the numerical rank of the symbol, and reflects the difficulty of storing it accurately, not the difficulty of computing it. In practice, we have found that typical values of C are still much smaller than the constants that arise in wavelet analysis, which are often plagued by a curse of dimensionality [13].

Wavelet matrices can sometimes be reduced in size to a mere $O(1)$ too, with controlled accuracy. To our knowledge this observation has not been reported in the literature yet, and goes to show that some care ought to be exercised before calling a method “optimal”. The particular smoothness properties of symbols that we leverage for their expansion is also hidden in the wavelet matrix, as *additional smoothness along the shifted diagonals*. The following result is elementary and we give it without proof.

Theorem 3. *Let $A \in \Psi^0$ as defined by (2), for $x \in \mathbb{R}$ and $\xi \in \mathbb{R}$. Let $\psi_{j,k}$ be an orthonormal wavelet basis of $L^2(\mathbb{R})$ of class C^∞ , and with an infinite number of vanishing moments. Then for each j , and each $\Delta k = k - k'$, there exists a function $f_{j,\Delta k} \in C^\infty(\mathbb{R})$ with smoothness constants independent of j , such that*

$$\langle \psi_{j,k}, A\psi_{j,k'} \rangle = f_{j,\Delta k}(2^{-j}k).$$

We would like to mention that similar ideas of smoothness along the diagonal have appeared in the context of seismic imaging, for the diagonal fitting of the so-called normal operator in a curvelet frame [26, 10]. In addition, the construction of second-generation bandlets for image processing is based on a similar phenomenon of smoothness along edges for the unitary recombination of MRA




Figure 2: Left: the standard 512-by-512 wavelet matrix of the differential operator considered in Figure 1, truncated to elements greater than 10^{-5} (white). Right: the 65-by-15 interaction matrix of DSC, for the same operator and a comparable accuracy, using a hierarchical splines expansion in ξ . The scale is the same for both pictures. (For visualization, an interpolated version of the matrix on the right is in Figure 1, top-right). Notice that the DSC matrix can be further compressed by a singular value decomposition, and in this example has numerical rank equal to 3, for a singular value cutoff at 10^{-5} . For values of N greater than 512, the wavelet matrix would increase in size in a manner directly proportional to N , while the DSC matrix would grow in size like $\log N$.

wavelet coefficients [40]. We believe that this last “alpertization” step could be of great interest in numerical analysis.

Theorem 3 hinges on the assumption of symbols in S^m , which is not met in the more general context of Calderón-Zygmund operators (CZO) considered by Meyer, Beylkin, Coifman, and Rokhlin. The class of CZO has been likened to a limited-smoothness equivalent to symbols of type $(1, 1)$ and order 0, i.e., symbols that obey

$$|\partial_\xi^\alpha \partial_x^\beta a(x, \xi)| \leq C_{\alpha, \beta} \langle \xi \rangle^{-|\alpha| + |\beta|}.$$

Symbols of type $(1, 0)$ and order 0 obeying (2) are a special case of this. Wavelet matrices of operators in the $(1, 1)$ class are almost diagonal⁴, but there is no smoothness along the shifted diagonals as in Theorem 3. So while the result in [3] is sharp, namely not much else than wavelet sparsity can be expected for CZO, we may question whether the generality of the CZO class is truly needed for applications to partial differential equations. The authors are unaware of a linear PDE setup involving symbols in the $(1, 1)$ class that would not also belong to the $(1, 0)$ class.

⁴Their standard wavelet matrix has at most $O(j)$ large elements per row and column at scale j —or frequency $O(2^j)$ —after which the matrix elements decay sufficiently fast below a preset threshold. L^2 boundedness would follow if there were $O(1)$ large elements per row and column, but $O(j)$ does not suffice for that, which testifies to the fact that operators of type $(1, 1)$ are not in general L^2 bounded. The reason for this $O(j)$ number is that an operator with a $(1, 1)$ symbol does not preserve vanishing moments of a wavelet—not even approximately. Such operators may turn an oscillatory wavelet at any scale j into a non-oscillating bump, which then requires wavelets at all the coarser scales for its expansion.

1.6 Related work

The idea of writing pseudodifferential symbols in separated form to formulate various one-way approximations to the variable-coefficients Helmholtz equation has long been a tradition in seismic imaging. This almost invariably involves a high-frequency approximation of some kind. Some influential work includes the phase screen method by Fisk and McCartor [19], and the generalized screen expansion of Le Rousseau and de Hoop [35]. This last reference discusses fast application of pseudodifferential operators in separated form using the FFT, and it is likely not the only reference to make this simple observation. A modern treatment of leading-order pseudodifferential approximations to one-way wave equations is in [48].

Expansions of principal symbols $a_0(x, \xi/|\xi|)$ (homogeneous of degree 0 in ξ) in spherical harmonics in ξ is a useful tool in the theory of pseudodifferential operators [49], and has also been used for fast computations by Bao and Symes in [1]. For computation of pseudodifferential operators, see also the work by Lamoureux, Margrave, and Gibson [34].

Symbol factorization has also served to design ILU preconditioners for the Helmholtz equation in [21] by Gander and Nataf. The notion of symbol is identical to that of generating function for Toeplitz or quasi-Toeplitz matrices: the algorithmic implications of approximating generating functions for preconditioning Toeplitz matrices are reported in [43]. An application of generating functions of Toeplitz matrices to the analysis of multigrid methods is in [31].

In the numerical analysis community, separation of operator kernels and other high-dimensional functions is becoming an important topic. Beylkin and Mohlenkamp proposed an alternated least-squares algorithm for computing separated expansions of tensors in [4, 5], propose to compute functions of operators in this representation, and apply these ideas to solving the multiparticle Schrödinger equation in [6], with Perez.

A different, competing approach to compressing operators is the “partitioned separated” method that consists in isolating off-diagonal squares of the kernel $K(x, y)$, and approximating each of them by a low-rank matrix. This also calls for an adapted notion of calculus, e.g., for composing and inverting operators. The first reference to this algorithmic framework is probably the partitioned SVD method described in [32]. More recently, these ideas have been extensively developed under the name H-matrix, for hierarchical matrix; see [9, 24] and <http://www.hlib.org>.

Separation ideas, with an adapted notion of operator calculus, have also been suggested for solving the wave equation; two examples are [7] and [14].

Exact operator square-roots — up to numerical errors — have in some contexts already been considered in the literature. See [18] for an example of Helmholtz operator with a quadratic profile, and [36] for a spectral approach that leverages sparsity, also for the Helmholtz operator.

2 Discrete Symbol Calculus: Representations

The two central questions of discrete symbol calculus are:

- Given an operator A , how to represent its symbol $a(x, \xi)$ efficiently?
- How to perform the basic operations of the pseudodifferential symbol calculus based on this representation? These operations include sum, product, adjoint, inversion, square root, inverse square root, and, in some cases, the exponential.

These two questions are mostly disjoint; we answer the first question in this section, and the second question in Section 3.

Let us write expansions of the form (5). Since $e_\lambda(x) = e^{2\pi i x \cdot \lambda}$ with $x \in [0, 1]^2$, we define the ξ -normalized x -Fourier coefficients of $a(x, \xi)$ as

$$a_\lambda(\xi) := \langle \xi \rangle^{-d_a} \int_{[0,1]^2} e^{-2\pi i x \cdot \lambda} a(x, \xi) dx, \quad \lambda \in \mathbb{Z}^2. \quad (13)$$

Note that the factor $\langle \xi \rangle^{-d_a}$ removes the growth or decay for large $|\xi|$. Clearly,

$$a(x, \xi) = \sum_{\lambda} e_\lambda(x) a_\lambda(\xi) \langle \xi \rangle^{d_a}. \quad (14)$$

In the case when $a(\cdot, \xi)$ is essentially bandlimited with band B_x , i.e., $a_\lambda(\xi)$ is supported inside the square $(-B_x, B_x)^2$ in the λ -frequency domain, then the integral in (13) can be approximated accurately by a uniform quadrature on the points $x_p = p/(2B_x)$, with $0 \leq p_1, p_2 < 2B_x$. This grid is called X in the sequel.

The problem is now reduced to finding an adequate approximation $\tilde{a}_\lambda(\xi)$ for $a_\lambda(\xi)$, either valid in the whole plane $\xi \in \mathbb{R}^2$, or in a large square $\xi \in [-N, N]^2$. Once this is done, then

$$\tilde{a}(x, \xi) := \sum_{\lambda \in (-B_x, B_x)^2} e_\lambda(x) \tilde{a}_\lambda(\xi) \langle \xi \rangle^{d_a}$$

is the desired approximation.

2.1 Rational Chebyshev interpolant

For symbols in the class (3), the function $a_\lambda(\xi)$ for each λ is smooth in angle $\arg \xi$, and polyhomogeneous in radius $|\xi|$. This means that $a_\lambda(\xi)$ is for $|\xi|$ large a polynomial of $1/|\xi|$ along each radial line through the origin, and is otherwise smooth (except possibly near the origin).

One idea for efficiently expanding such functions is to map the half line $|\xi| \in [0, \infty)$ to the interval $[-1, 1]$ by a rational function, and expand the result in Chebyshev polynomials. Put $\xi = (\theta, r)$, and $\mu = (m, n)$. Let

$$g_\mu(\xi) = e^{im\theta} TL_n(r),$$

where TL_n are the rational Chebyshev functions [8], defined from the Chebyshev polynomials of the first kind T_n as

$$TL_n(r) = T_n(A_L^{-1}(r)),$$

by means of the algebraic map

$$s \mapsto r = A_L(s) = L \frac{1+s}{1-s}, \quad r \mapsto s = A_L^{-1}(r) = \frac{r-L}{r+L}.$$

The parameter L is typically on the order of 1. The proposed expansion then takes the form

$$a_\lambda(\xi) = \sum_{\mu} a_{\lambda, \mu} g_\mu(\xi),$$

also denoted $\tilde{h}_{a, \lambda}(\xi)$ in the sequel when the sum is truncated, and where

$$a_{\lambda, \mu} = \frac{1}{2\pi} \int_{-1}^1 \int_0^{2\pi} a_\lambda((\theta, A_L(s))) e^{-im\theta} T_n(s) \frac{d\theta ds}{\sqrt{1-s^2}}.$$

For properly bandlimited functions, such integrals can be evaluated exactly using the right quadrature points: uniform in $\theta \in [0, 2\pi]$, and Chebyshev points in s . The corresponding points

in r are the image of the Chebyshev points under the algebraic map. The resulting grid in the ξ plane can be described as follows. Let $q = (q_\theta, q_r)$ be a couple of integers such that $0 \leq q_\theta < N_\theta$ and $0 \leq q_r < N_r$; we have in polar coordinates

$$\xi_q = \left(2\pi \frac{q_\theta}{N_\theta}, -\cos \left(\frac{2(A_L(q_r) - 1)}{2N_r} \right) \right).$$

We call this grid $\{\xi_q\} = \Omega$. Passing from the values $a_\lambda(\xi_q)$ to $a_{\lambda,\mu}$ and vice-versa can be done using the fast Fourier transform. Of course, $\tilde{a}_\lambda(\xi)$ is nothing but an interpolant of $a_\lambda(\xi)$ at the points ξ_q .

In the remainder of this section, we present the proof of Theorem 1, which contains the convergence rates of the truncated sums over λ and μ . The argument hinges on the following L^2 boundedness result, which is a simple modification of standard results in \mathbb{R}^d , see [46]. It is not necessary to restrict $d = 2$ for this lemma.

Lemma 1. *Let $a(x, \xi) \in C^{d'}([0, 1]^d, \ell_\infty(\mathbb{Z}^d))$, where $d' = d + 1$ if d is odd, or $d + 2$ if d is even. Then the operator A defined by (6) extends to a bounded operator on $L^2([0, 1]^d)$, with*

$$\|A\|_{L^2} \leq C \cdot \|(1 + (-\Delta_x)^{d'/2})a(x, \xi)\|_{L^\infty([0, 1]^d, \ell_\infty(\mathbb{Z}^d))}.$$

The proof of this lemma is in the Appendix.

Proof of Theorem 1. Consider $s = A_L^{-1}(r) \in [-1, 1)$ where A_L and its inverse were defined above. Expanding $a(x, (\theta, r))$ in rational Chebyshev functions $TL_n(r)$ is equivalent to expanding $f(s) \equiv a(x, (\theta, A_L(s)))$ in Chebyshev polynomials $T_n(s)$. Obviously,

$$f \circ A_L^{-1} \in C^\infty([0, \infty)) \quad \Leftrightarrow \quad f \in C^\infty([-1, 1)).$$

It is furthermore assumed that $a(x, \xi)$ is in the classical class with tempered growth of the polyhomogeneous components; this condition implies that the smoothness constants of $f(s) = a(x, (\theta, A_L(s)))$ are uniform as $s \rightarrow 1$, i.e., for all $n \geq 0$,

$$\exists C_n : |f^{(n)}(s)| \leq C_n, \quad s \in [-1, 1],$$

or simply, $f \in C^\infty([-1, 1])$. In order to see why that is the case, consider a cutoff function $\chi(r)$ equal to 1 for $r \geq 2$, zero for $0 \leq r \leq 1$, and C^∞ increasing in between. Traditionally, the meaning of (3) is that there exists a sequence $\varepsilon_j > 0$, defining cutoffs $\chi(r\varepsilon_j)$ such that

$$a(x, (r, \theta)) - \sum_{j \geq 0} a_j(x, \theta) r^{-j} \chi(r\varepsilon_j) \in S_{\text{cl}}^{-m}, \quad \forall m > 0.$$

A remainder in $S_{\text{cl}}^{-\infty} \equiv \bigcup_{m \geq 0} S_{\text{cl}}^{-m}$ is called smoothing. As long as the choice of cutoffs ensures convergence, the determination of $a(x, \xi)$ modulo $S^{-\infty}$ does not depend on this choice. (Indeed, if there existed an order- k discrepancy between the sums with $\chi(r\varepsilon_j)$ or $\chi(r\delta_j)$, with k finite, it would need to come from some of the terms $a_j r^{-j} (\chi(r\varepsilon_j) - \chi(r\delta_j))$ for $j \leq k$. But each of these terms is of order $-\infty$.)

Because of condition (7), it is easy to check that the particular choice $\varepsilon_j = 1/(2R)$ suffices for convergence of the sum over j to a symbol in S^0 . As mentioned above, changing the ε_j only affects the smoothing remainder, so we may focus on $\varepsilon_j = 1/(2R)$.

After changing variables, we get

$$f(s) = a(x, (\theta, A_L(s))) = \sum_{j \geq 0} a_j(x, \theta) L^{-j} \left(\frac{1-s}{1+s} \right)^j \chi \left(\frac{A_L(s)}{2R} \right) + r(s),$$

where the smoothing remainder $r(s)$ obeys

$$|r^{(n)}(s)| \leq C_{n,M}(1-s)^M, \quad \forall M \geq 0,$$

hence, in particular when $M = 0$, has uniform smoothness constants as $s \rightarrow 1$. It suffices therefore to show that the sum over $j \geq 0$ can be rewritten as a Taylor expansion for $f(s) - r(s)$, convergent in some neighborhood of $s = 1$.

Let $z = 1 - s$. Without loss of generality, assume that $R \geq 2L$, otherwise increase R to $2L$. The cutoff factor $\chi\left(\frac{A_L(1-z)}{2R}\right)$ equals 1 as long as $0 \leq z \leq \frac{L}{4R}$. In that range,

$$f(1-z) - r(1-z) = \sum_{j \geq 0} a_j(x, \theta) L^{-j} \frac{z^j}{(2-z)^j}.$$

By making use of the binomial expansion

$$\frac{z^j}{(2-z)^j} = \sum_{m \geq 0} \left(\frac{z}{2}\right)^{j+m} \binom{j+m-1}{j-1}, \quad \text{if } j \geq 1,$$

and the new index $k = j + m$, we obtain the Taylor expansion about $z = 0$:

$$f(1-z) - r(1-z) = a_0(x, \theta) + \sum_{k \geq 0} \left(\frac{z}{2}\right)^k \sum_{1 \leq j \leq k} \frac{a_j(x, \theta)}{L^j} \binom{k-1}{j-1}.$$

To check convergence, notice that $\binom{k-1}{j-1} \leq \sum_{n=0}^{k-1} \binom{k-1}{n} = 2^{k-1}$, combine this with (7), and obtain

$$\begin{aligned} 2^{-k} \sum_{1 \leq j \leq k} \frac{a_j(x, \theta)}{L^j} \binom{k-1}{j-1} &\leq \frac{Q_{00}}{2} \sum_{1 \leq j \leq k} \left(\frac{R}{L}\right)^j \\ &\leq \frac{Q_{00}}{2} \frac{1}{1-L/R} \left(\frac{R}{L}\right)^k \end{aligned}$$

We assumed earlier that $z \in [0, L/(4R)]$: this condition manifestly suffices for convergence of the sum over k . This shows that $f \in C^\infty([-1, 1])$; the very same reasoning with $Q_{\alpha\beta}$ in place of Q_{00} also shows that any derivative $\partial_x^\alpha \partial_\theta^\beta f(s) \in C^\infty([-1, 1])$.

The Chebyshev expansion of $f(s)$ is the Fourier-cosine series of $f(\cos \phi)$, with $\phi \in [0, \pi]$; the previous reasoning shows that $f(\cos \phi) \in C^\infty([0, \infty])$. The same is true for any (x, θ) derivatives of $f(\cos \phi)$.

Hence $a(x, (A_L(\cos \phi), \theta))$ is a C^∞ function, periodic in all its variables. The proposed expansion scheme is simply:

- A Fourier series in $x \in [0, 1]^2$;
- A Fourier series in $\theta \in [0, 2\pi]$;
- A Fourier-cosine series in $\phi \in [0, \pi]$.

An approximant with at most M terms can then be defined by keeping $\lfloor M^{1/4} \rfloor$ Fourier coefficients per direction. It is well-known that Fourier and Fourier-cosine series of a C^∞ , periodic function

converge super-algebraically in the L^∞ norm, and that the same is true for any derivative of the function as well. Therefore if a_M is this M -term approximant, we have

$$\sup_{x,\theta,\phi} |\partial_x^\beta (a - \tilde{a})(x, (A_L(\cos \phi), \theta))| \leq C_{\beta,M} \cdot M^{-\infty}, \quad \forall \text{ multi-index } \beta$$

We now invoke Lemma 1 with $a - a_M$ in place of a , choose $M = O(\varepsilon^{-1/\infty})$ with the right constants, and conclude. \square

It is interesting to observe what goes wrong when condition (7) is not satisfied. For instance, if the growth of the a_j is fast enough in (3), then it may be possible to choose the cutoffs $\chi(\varepsilon_j|\xi|)$ such that the sum over j replicates a fractional negative power of $|\xi|$, like $|\xi|^{-1/2}$, and in such a way that the resulting symbol is still in the class defined by (2). A symbol with this kind of decay at infinity would not be mapped onto a C^∞ function of s inside $[-1, 1]$ by the algebraic change of variables A_L , and the Chebyshev expansion in s would not converge spectrally. This kind of pathology is generally avoided in the literature on pseudodifferential operators by assuming that the order of the compound symbol $a(x, \xi)$ is the same as that of the principal symbol, i.e., the leading-order contribution $a_0(x, \arg \xi)$.

Finally, note that the obvious generalization of the complex exponentials in $\arg \xi$ to higher-dimensional settings would be spherical harmonics, as advocated in [1]. The radial expansion scheme would remain unchanged.

2.2 Hierarchical spline interpolant

An alternative representation is to use a hierarchical spline construction in the ξ plane. We define $\tilde{a}_\lambda(\xi)$ to be an interpolant of the ξ -normalized x -Fourier coefficients $a_\lambda(\xi)$ as follows. The interpolant is defined in the square $\xi \in [-N, N]^2$ for some large N . Pick a number B_ξ , independent of N , that plays the role of coarse-scale bandwidth. In practice it is taken comparable to B_x .

- Define $D_0 = (-B_\xi, B_\xi)^2$. For each $\xi \in D_0$, $\tilde{a}_\lambda(\xi) := a_\lambda(\xi)$.
- For each $\ell = 1, 2, \dots, L = \log_3(N/B_\xi)$, define $D_\ell = (-3^\ell B_\xi, 3^\ell B_\xi)^2 \setminus D_{\ell-1}$. D_ℓ is further partitioned into eight blocks:

$$D_\ell = \bigcup_{i=1}^8 D_{\ell,i},$$

where each block $D_{\ell,i}$ is of size $2 \cdot 3^{\ell-1} B_\xi \times 2 \cdot 3^{\ell-1} B_\xi$. Within each block $D_{\ell,i}$, sample $a_\lambda(\xi)$ with a Cartesian grid $G_{\ell,i}$ of a fixed size. The restriction of $\tilde{a}_\lambda(\xi)$ in $D_{\ell,i}$ is defined to be the spline interpolant of $a_\lambda(\xi)$ on the grid $G_{\ell,i}$.

We emphasize that the number of samples used in each grid $G_{\ell,i}$ is fixed independent of the level ℓ . The reason for this choice is that the function $a_\lambda(\xi)$ gains smoothness as ξ grows to infinity. In practice, we set $G_{\ell,i}$ to be a 4×4 or 5×5 Cartesian grid and use cubic spline interpolation.

Let us summarize the construction of the representation $\tilde{a}(x, \xi) = \sum_\lambda e_\lambda(x) \tilde{a}_\lambda(\xi) \langle \xi \rangle^{d_a}$. As before, fix a parameter B_x that governs the bandwidth in x and define

$$X = \left\{ \left(\frac{p_1}{2B_x}, \frac{p_2}{2B_x} \right), 0 \leq p_1, p_2 < 2B_x \right\} \quad \text{and} \quad \Omega = D_0 \bigcup \left(\bigcup_{\ell,i} G_{\ell,i} \right).$$

The construction of the expansion of $a(x, \xi)$ takes the following steps

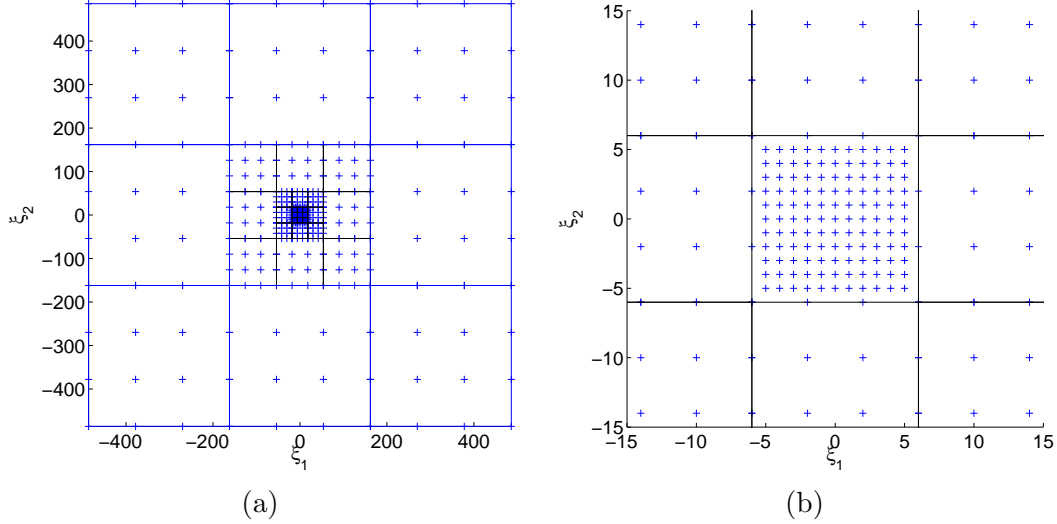


Figure 3: Hierarchical spline construction. Here $B_\xi = 6$, $L = 4$, and $N = 486$. The grid $G_{\ell,i}$ is of size 4×4 . The grid points are shown with “+” sign. (a) The whole grid. (b) The center of the grid.

- Sample $a(x, \xi)$ for all pairs of (x, ξ) with $x \in X$ and $\xi \in \Omega$.
- For a fixed $\xi \in \Omega$, use the fast Fourier transform to compute $a_\lambda(\xi)$ for all $\lambda \in (-B_x, B_x)^2$.
- For each λ , construct the interpolant $\tilde{a}_\lambda(\xi)$ from the values of $a_\lambda(\xi)$.

Let us study the complexity of this construction procedure. The number of samples in X is bounded by $4B_x^2$, considered a constant with respect to N . As we use a constant number of samples for each level $j = 1, 2, \dots, L = \log_3(N/B_\xi)$, the number of samples in Ω is of order $O(\log N)$. Therefore, the total number of samples is still of order $O(\log N)$. Similarly, since the construction of a fixed size spline interpolant requires only a fixed number of steps, the construction of the interpolants $\{\tilde{a}_\lambda(\xi)\}$ takes only $O(\log N)$ steps as well. Finally, we would like to remark that, due to the locality of the spline, the evaluation of $\tilde{a}_\lambda(\xi)$ for any fixed λ and ξ requires only a constant number of steps.

We now expand on the convergence properties of the spline interpolant.

Proof of Theorem 2. If the number of control points per square $D_{j,i}$ is K^2 instead of 16 or 25 as we advocated above, the spline interpolant becomes arbitrarily accurate. The spacing between two control points at level j is $O(3^j/K)$. With p be the order of the spline scheme — we took $p = 3$ earlier — it is standard polynomial interpolation theory that

$$\sup_{\xi \in D_{j,i}} |\tilde{a}_\lambda(\xi) - a_\lambda(\xi)| \leq C_{a,\lambda,p} \cdot \left(\frac{3^j}{K}\right)^{p+1} \cdot \sup_{|\alpha|=p+1} \|\partial_\xi^\alpha a_\lambda\|_{L^\infty(D_{j,i})}.$$

The symbol estimate (2) guarantees that the last factor is bounded by $C \cdot \sup_{\xi \in D_{j,i}} \langle \xi \rangle^{-p-1}$. Each square $D_{j,i}$, for fixed j , is at a distance $O(3^j)$ from the origin, hence $\sup_{\xi \in D_{j,i}} \langle \xi \rangle^{-p-1} = O(3^{-j(p+1)})$. This results in

$$\sup_{\xi \in D_{j,i}} |\tilde{h}_{a,\lambda}(\xi) - h_{a,\lambda}(\xi)| \leq C_{a,\lambda,p} \cdot K^{-p-1}.$$

This estimate is uniform over $D_{j,i}$, hence also over $\xi \in [-N, N]^2$. As argued earlier, it is achieved by using $O(K^2 \log N)$ spline control points. If we factor in the error of expanding the symbol in the x variable using $4B^2$ spatial points, for a total of $M = O(B^2 K^2 \log N)$ points, we have the compound estimate

$$\sup_{x \in [0,1]^2} \sup_{\xi \in [-N,N]^2} |a(x, \xi) - \tilde{a}(x, \xi)| \leq C \cdot (B^{-\infty} + K^{-p-1}).$$

The same estimate holds for the partial derivatives of $a - \tilde{a}$ in x .

Functions to which the operator defined by $\tilde{a}(x, \xi)$ is applied need to be bandlimited to $[-N, N]^2$, i.e., $\hat{f}(\xi) = 0$ for $\xi \notin [-N, N]^2$, or better yet $f = P_N f$. For those functions, the symbol \tilde{a} can be extended by a outside of $[-N, N]^2$, Lemma 1 can be applied to the difference $A - \tilde{A}$, and we obtain

$$\|(A - \tilde{A})f\|_{L^2} \leq C \cdot (B^{-\infty} + K^{-p-1}) \cdot \|f\|_{L^2}.$$

The leading factors of $\|f\|_{L^2}$ in the right-hand side can be made less than ε if we choose $B = O(\varepsilon^{-1/\infty})$ and $K = O(\varepsilon^{-1/(p+1)})$, with adequate constants. The corresponding number of points in x and ξ is therefore $M = O(\varepsilon^{-2/(p+1)} \cdot \log N)$. □

3 Discrete Symbol Calculus: Operations

Let A and B be two operators with symbols $a(x, \xi)$ and $b(x, \xi)$. Suppose that we have already generated their expansions

$$a(x, \xi) \approx \tilde{a}(x, \xi) = \sum_{\lambda} e_{\lambda}(x) \tilde{a}_{\lambda}(\xi) \langle \xi \rangle^{d_a} \quad \text{and} \quad b(x, \xi) \approx \tilde{b}(x, \xi) = \sum_{\lambda} e_{\lambda}(x) \tilde{b}_{\lambda}(\xi) \langle \xi \rangle^{d_b}.$$

Here d_a and d_b are the orders of $a(x, \xi)$ and $b(x, \xi)$, respectively. It is understood that the sum over λ is restricted to $(-B_x, B_x)^2$, that $a_{\lambda}(\xi)$ are approximated with $\tilde{a}_{\lambda}(\xi)$ by either method described earlier, and that we will not keep track of which particular method is used in the notations. Let us now consider the basic operations of the calculus of discrete symbols.

Scaling $C = \alpha A$. For the symbols, we have $c(x, \xi) = \alpha a(x, \xi)$. In terms of the Fourier coefficients,

$$c_{\lambda}(\xi) \langle \xi \rangle^{d_a} = \alpha a_{\lambda}(\xi) \langle \xi \rangle^{d_a}.$$

Therefore, we set $d_c = d_a$ and take the approximant $\tilde{c}_{\lambda}(\xi)$ to be

$$\tilde{c}_{\lambda}(\xi) := \alpha \cdot \tilde{a}_{\lambda}(\xi).$$

Sum $C = A + B$. For the symbols, we have $c(x, \xi) = a(x, \xi) + b(x, \xi)$. In terms of the Fourier coefficients,

$$c_{\lambda}(\xi) \langle \xi \rangle^{d_c} = a_{\lambda}(\xi) \langle \xi \rangle^{d_a} + b_{\lambda}(\xi) \langle \xi \rangle^{d_b}.$$

Therefore, it is natural to set $d_c = \max(d_a, d_b)$ and $\tilde{c}_{\lambda}(\xi)$ to be the interpolant with values

$$\left(\tilde{a}_{\lambda}(\xi) \langle \xi \rangle^{d_a} + \tilde{b}_{\lambda}(\xi) \langle \xi \rangle^{d_b} \right) \langle \xi \rangle^{-d_c}$$

for $\xi \in \Omega$. Here Ω is either the Chebyshev points grid, or the hierarchical spline grid defined earlier.

Product $C = AB$. For the symbols, we have

$$c(x, \xi) = a(x, \xi) \# b(x, \xi) = \sum_{\eta} \int e^{-2\pi i(x-y)(\xi-\eta)} a(x, \eta) b(y, \xi) dy.$$

In terms of the Fourier coefficients,

$$c_{\lambda}(\xi) \langle \xi \rangle^{d_c} = \sum_{k+l=\lambda} a_k(\xi+l) \langle \xi+l \rangle^{d_a} b_l(\xi) \langle \xi \rangle^{d_b}.$$

Therefore, $d_c = d_a + d_b$ and $\tilde{c}_{\lambda}(\xi)$ is taken to be the interpolant with values

$$\left(\sum_{k+l=\lambda} \tilde{a}_k(\xi+l) \langle \xi+l \rangle^{d_a} \tilde{b}_l(\xi) \langle \xi \rangle^{d_b} \right) \langle \xi \rangle^{-d_c}$$

at $\xi \in \Omega$.

Transpose $C = A^*$. It is straightforward to derive the formula of its symbol:

$$c(x, \xi) = \sum_{\eta} \int e^{-2\pi i(x-y)(\xi-\eta)} \overline{a(y, \eta)} dy.$$

In terms of the Fourier coefficients,

$$c_{\lambda}(\xi) \langle \xi \rangle^{d_c} = \overline{a_{-\lambda}(\xi+\lambda)} \langle \xi+\lambda \rangle^{d_a}.$$

Therefore, $d_c = d_a$ and $\tilde{c}_{\lambda}(\xi)$ is the interpolant that takes the values

$$\left(\overline{\tilde{a}_{-\lambda}(\xi+\lambda)} \langle \xi+\lambda \rangle^{d_a} \right) \langle \xi \rangle^{-d_c}$$

at $\xi \in \Omega$.

Inverse $C = A^{-1}$ where A is symmetric positive definite. We first pick a constant α such that $\alpha|a(x, \xi)| \ll 1$ for $\xi \in (-N, N)^2$. Since the order of $a(x, \xi)$ is d_a , $\alpha \approx O(1/N^{d_a})$. In the following iteration, we first invert αA and then scale the result by α to get C .

- $X_0 = I$.
- For $k = 0, 1, 2, \dots$, repeat $X_{k+1} = 2X_k - X_k(\alpha A)X_k$ until convergence.
- Set $C = \alpha X_k$.

This iteration is called the Schulz iteration⁵, and is quoted in [4]. It can be seen as a modified Newton iteration for finding the nontrivial zero of $f(X) = XAX - X$, where the gradient of f is approximated by the identity.

As this algorithm only utilizes the addition and the product of the operators, all of the computation can be carried out via discrete symbol calculus. Since $\alpha \approx O(1/N^{d_a})$, the smallest eigenvalue of αA can be as small as $O(1/N^{d_a})$ where the constant depends on the smallest eigenvalue of A . For a given accuracy ε , it is not difficult to heuristically show that this algorithm converges after $O(\log N + \log(1/\varepsilon))$ iterations. The constant in this estimate is proportional to d_a , i.e., proportional to the logarithm of the condition number of A .

⁵We thank a referee for the original reference [41] to Schulz, and for pointing out that it is also sometimes called the Hotelling-Bodewig iteration.

Square root and inverse square root Put $C = A^{1/2}$ and $D = A^{-1/2}$ where A is symmetric positive definite. Here, we again choose a constant α such that $\alpha|a(x, \xi)| \ll 1$ for $\xi \in (-N, N)^2$. This also implies that $\alpha \approx O(1/N^{d_a})$. In the following iteration, the Schulz-Higham iteration [27, 28, 30, 41] is used to compute the square root and the inverse square root of αA and these operators are scaled appropriately.

- $Y_0 = \alpha A$ and $Z_0 = I$.
- For $k = 0, 1, 2, \dots$, repeat $Y_{k+1} = \frac{1}{2}Y_k(3I - Z_k Y_k)$ and $Z_{k+1} = \frac{1}{2}(3I - Z_k Y_k)Z_k$ until convergence.
- Set $C = \alpha^{-1/2}Y_k$ and $D = \alpha^{1/2}Z_k$.

In a similar way to the iteration used for computing the inverse, the Schulz-Higham iteration is similar to the iteration for computing the inverse in that it uses only additions and products. Therefore, all of the computation can be performed via discrete symbol calculus. A similar analysis show that, for any fixed accuracy ε , the number of iterations required by the Schulz-Higham iteration is of order $O(\log N + \log(1/\varepsilon))$, as for the inverse.

Exponential $C = e^{\alpha A}$. In general, the exponential of an elliptic pseudodifferential operator is not necessarily a pseudodifferential operator itself. However, if the data is restricted to $\xi \in (-N, N)^2$ and $\alpha = O(1/N^{d_a})$, the exponential operator behaves almost like a pseudodifferential operator in this range of frequencies⁶. In Section 4.4, we will give an example where such an exponential operator plays an important role.

We construct C using the following scaling-and-squaring steps [39]:

- Pick δ sufficient small so that $\alpha/\delta = 2^K$ for an integer K .
- Construct an approximation Y_0 for $e^{\delta A}$. One possible choice is the 4th order Taylor expansion: $Y_0 = I + \delta A + \frac{(\delta A)^2}{2!} + \frac{(\delta A)^3}{3!} + \frac{(\delta A)^4}{4!}$. Since δ is sufficient small, Y_0 is quite accurate.
- For $k = 0, 1, 2, \dots, K - 1$, repeat $Y_{k+1} = Y_k Y_k$.
- Set $C = Y_K$.

This iteration for computing the exponential again uses only the addition and product operations and, therefore, all the steps can be carried out at the symbol level using discrete symbol calculus. The number of steps K is usually quite small, as the constant α itself is of order $O(1/N^{d_a})$.

Moyal transform Pseudodifferential operators are sometimes defined by means of their Weyl symbol a_W , as

$$Af(x) = \sum_{\xi \in \mathbb{Z}^d} \int_{[0,1]^d} a_W\left(\frac{1}{2}(x+y), \xi\right) e^{2\pi i(x-y)\xi} f(y) dy,$$

when $\xi \in \mathbb{Z}^d$, otherwise if $\xi \in \mathbb{R}^d$, replace the sum over ξ by an integral. It is a more symmetric formulation that may be preferred in some contexts. The other, usual formulation we have used throughout this paper is called the Kohn-Nirenberg correspondence. The relationship between the

⁶Note that another case in which the exponential remains pseudodifferential is when the spectrum of A is real and negative, regardless of the size of α .

two methods of “quantization”, i.e., passing from a symbol to an operator, is the so-called Moyal transform. The book [20] gives the recipe:

$$a_W(x, \xi) = (Ma)(x, \xi) = 2^n \sum_{\eta \in \mathbb{Z}^d} \int e^{4\pi i(x-y) \cdot (\xi-\eta)} a(y, \eta) dy,$$

and conversely

$$a(x, \xi) = (M^{-1}a_W)(x, \xi) = 2^n \sum_{\eta \in \mathbb{Z}^d} \int e^{-4\pi i(x-y) \cdot (\xi-\eta)} a_W(y, \eta) dy.$$

These operations are algorithmically very similar to transposition. It is interesting to notice that transposition is a mere conjugation in the Weyl domain: $a^* = M^{-1}(\overline{Ma})$. We also have the curious property that

$$\widehat{Ma}(p, q) = e^{-\pi i p q} \widehat{a}(p, q)$$

where the hat denotes Fourier transform in both variables.

Applying the operator The last operation that we discuss is how to apply the operator to a given input function. Suppose $u(x)$ is sampled on a grid $x = (p_1/P, p_2/P)$ with $0 \leq p_1, p_2 < P$ and $P/2 < N$. Our goal is to compute $(Au)(x)$ on the same grid. Using the definition of the pseudo-differential symbol and the expansion of $a(x, \xi)$, we have

$$\begin{aligned} (Au)(x) &= \sum_{\xi} e^{2\pi i x \xi} a(x, \xi) \widehat{u}(\xi) \\ &\approx \sum_{\xi} e^{2\pi i x \xi} \sum_{\lambda} e_{\lambda}(x) \tilde{a}_{\lambda}(\xi) \langle \xi \rangle^{d_a} \widehat{u}(\xi) \\ &= \sum_{\lambda} e_{\lambda}(x) \left(\sum_{\xi} e^{2\pi i x \xi} \left(\tilde{a}_{\lambda}(\xi) \langle \xi \rangle^{d_a} \widehat{u}(\xi) \right) \right). \end{aligned}$$

Therefore, a straightforward way to compute Au is

- For each $\lambda \in (-B_x, B_x)^2$, sample $\tilde{h}_{a,\lambda}(\xi)$ for $\xi \in [-P/2, P/2]^2$.
- For each $\lambda \in (-B_x, B_x)^2$, form the product $\tilde{a}_{\lambda}(\xi) \langle \xi \rangle^{d_a} \widehat{u}(\xi)$ for $\xi \in [-P/2, P/2]^2$.
- For each $\lambda \in (-B_x, B_x)^2$, apply the fast Fourier transform to the result of the previous step.
- For each $\lambda \in (-B_x, B_x)^2$, multiply the result of the previous step with $e_{\lambda}(x)$. Finally, their sum gives $(Au)(x)$.

Let us estimate the complexity of this procedure. For each fixed λ , the number of operations is dominated by the complexity of the fast Fourier transform, which is $O(P^2 \log P)$. Since there is only a constant number of values for $\lambda \in (-B_x, B_x)^2$, the overall complexity is also $O(P^2 \log P)$.

In many cases, we need to calculate $(Au)(x)$ for many different functions $u(x)$. Though the above procedure is quite efficient, we can further reduce the number of the Fourier transforms required. The idea is to exploit the possible redundancy between the functions $\tilde{a}_{\lambda}(\xi)$ for different λ . We first use a rank-reduction procedure, such as QR factorization or singular value decomposition (SVD), to obtain a low-rank approximation

$$\tilde{a}_{\lambda}(\xi) \approx \sum_{t=1}^T u_{\lambda t} v_t(\xi) \tag{15}$$

where the number of terms T is often much smaller than the possible values of λ . We can then write

$$\begin{aligned} (Au)(x) &\approx \sum_{\lambda} e_{\lambda}(x) \sum_{\xi} e^{2\pi i x \xi} \sum_{t=1}^T u_{\lambda t} v_t(\xi) \langle \xi \rangle^{d_a} \hat{u}(\xi) \\ &= \sum_{t=1}^T \left(\sum_{\lambda} e_{\lambda}(x) u_{\lambda t} \right) \left(\sum_{\xi} e^{2\pi i x \xi} v_t(\xi) \langle \xi \rangle^{d_a} \hat{u}(\xi) \right). \end{aligned}$$

The improved version of applying $(Au)(x)$ then takes two steps. In the preprocessing step, we compute

- For each $\lambda \in (-B_x, B_x)^2$, sample $\tilde{a}_{\lambda}(\xi)$ for $\xi \in [-P/2, P/2]^2$.
- Construct the factorization $\tilde{a}_{\lambda}(\xi) \approx \sum_{t=1}^T u_{\lambda t} v_t(\xi)$.
- For each t , compute the function $\sum_{\lambda} e_{\lambda}(x) u_{\lambda t}$.

In the evaluation step, the following steps is carried out for an input function $u(x)$

- For each t , compute $v_t(\xi) \langle \xi \rangle^{d_a} \hat{u}(\xi)$.
- For each t , perform fast Fourier transform to the result of the previous step.
- For each t , multiply the result with $\sum_{\lambda} e_{\lambda}(x) u_{\lambda t}$. Their sum gives $(Au)(x)$.

4 Applications and Numerical Results

In this section, we provide several numerical examples to demonstrate the effectiveness of the discrete symbol calculus. In these numerical experiments, we use the hierarchical spline version of the discrete symbol calculus. Our implementation is written in Matlab and all the computational results are obtained on a desktop computer with a 2.8GHz CPU.

4.1 Basic operations

We first study the performance of the basic operations described in Section 3. In the following tests, we set $B_{\xi} = 6$, $L = 6$, and $N = B_{\xi} \times 3^L = 4374$. The number of samples in Ω is equal to 677. We consider the elliptic operator

$$Au := (I - \operatorname{div}(\alpha(x)\nabla))u.$$

The presence of the identity is unessential: it only makes inversion meaningful. For periodic boundary conditions, $\operatorname{div}(\alpha(x)\nabla)$ has a nonzero nullspace. It would be a very similar numerical task to remove this nullspace by instead prescribing the value of the symbol of A^{-1} to be zero at the origin in ξ . Even though the symbol of A^{-1} has a singularity at $\xi = 0$ for the continuous problem, the problem disappears when $\xi \in \mathbb{Z}^d$ away from the origin. As explained earlier, the complexity of discrete symbol calculus is only very mildly affected by the conditioning of A .

Example 1. The coefficient $\alpha(x)$ of this example is a simple sinusoid function given in Figure 4. We apply the discrete symbol calculus to the computation of the operators $C = AA$, $C = A^{-1}$, and $C = A^{1/2}$. The error is estimated by applying these operators to random noise test functions. For a given test function f , the errors are computed using:

- $\frac{\|Cf - A(Af)\|}{\|A(Af)\|}$ for $C = AA$,
- $\frac{\|A(Cf) - f\|}{\|f\|}$ for $C = A^{-1}$,
- $\frac{\|C(Cf) - Af\|}{\|Af\|}$ for $C = A^{1/2}$.




Figure 4: Coefficient $\alpha(x)$ of Example 1.

	Iter	Time(s)	Error
$C = AA$	-	3.66e+00	1.92e-05
$C = A^{-1}$	17	1.13e+02	2.34e-04
$C = A^{1/2}$	27	4.96e+02	4.01e-05

Table 1: Results of Example 1. The number of iterations, running time, and error of computing $C = AA$, $C = A^{-1}$, and $C = A^{1/2}$ with discrete symbol calculus.

We summarize in Table 1 the running time, the number of iterations, and the accuracy of these operations. Our algorithms produce good accuracy with a small number of sampling points in both x and ξ . The computation of the symbols of the inverse and the square root takes only a couple of minutes on a desktop computer, even for a large frequency domain $(-N, N)^2$ with $N = 4374$. Moreover, one can easily triple the value of N by adding one extra level in the hierarchical spline construction or by adding a few more radial quadrature points in the rational Chebyshev polynomial construction. In both cases, the running time and iteration count depend on N in a logarithmic way. This is in direct contrast with all other algorithms for constructing inverses and square roots of elliptic operators, where the complexity grows at least linearly with N , even with tools such as wavelets, hierarchical matrices, etc.

As we mentioned earlier, once A^{-1} is computed the computation of $A^{-1}f$ requires only a small number of Fourier transforms. Here, we compare our approach with the preconditioned conjugate gradient (PCG) algorithm, which is arguably one of the most efficient algorithms for the problem under consideration. The preconditioner we use is $M = I - \bar{\alpha}\Delta$ with $\bar{\alpha}$ taken to be the mean of α . In Table 2, we compare the running times of these two approaches. The function $f(x)$ is taken to be a random noise discretized on a uniform grid of size $P \times P$. In both approaches, the relative

P	DSC Time(s)	PCG Time(s)
128	5.00e-02	1.00e-01
256	1.90e-01	4.40e-01
512	9.50e-01	2.05e+00
1024	5.06e+00	1.46e+01

Table 2: Results of Example 1. The running times of computing $A^{-1}f$ using the DSC approach and the PCG algorithm for different problem sizes.

error is set at the order of $1e - 4$. Table 2 shows that two algorithm scale in the same way and the DSC approach is slightly faster.

Example 2. In this example, we set $\alpha(x)$ to be a random bandlimited function (see Figure 5). The running time, the number of iterations, and the error for each operation are reported in Table 3. A similar comparison with the PCG algorithm is given in Table 4.




Figure 5: Coefficient $\alpha(x)$ of Example 2.

	Iter	Time(s)	Error
$C = AA$	-	3.66e+00	1.73e-05
$C = A^{-1}$	16	1.05e+02	6.54e-04
$C = A^{1/2}$	27	4.96e+02	8.26e-05

Table 3: Results of Example 2. The number of iterations, running time, and error of computing $C = AA$, $C = A^{-1}$, and $C = A^{1/2}$ with discrete symbol calculus.

P	DSC Time(s)	PCG Time(s)
128	6.00e-02	1.00e-01
256	4.10e-01	4.40e-01
512	2.01e+00	2.05e+00
1024	9.59e+00	1.44e+01

Table 4: Results of Example 2. The running times of computing $A^{-1}f$ using the DSC approach and the PCG algorithm for different problem sizes.

From Tables 1 and 3, we observe that the number of iterations for the inverse and square root operator remain rather independent of the function $a(x, \xi)$.

4.2 Preconditioner

As we mentioned in the Introduction, an important application of the discrete symbol calculus is to precondition the inhomogeneous Helmholtz equation:

$$Lu := \left(-\Delta - \frac{\omega^2}{c^2(x)} \right) u = f$$

where the sound speed $c(x)$ is smooth and periodic in x . We consider the solution of the preconditioned system

$$M^{-1}Lu = M^{-1}f$$

with the so-called complex-shifted Laplace preconditioner [16], of which we consider two variants,

$$M_1 := -\Delta + \frac{\omega^2}{c^2(x)} \quad \text{and} \quad M_2 := -\Delta + (1 + i) \cdot \frac{\omega^2}{c^2(x)}.$$

For each preconditioner M_j with $j = 1, 2$, we use the discrete symbol calculus to compute the symbol of M_j^{-1} . As we mentioned earlier, applying M_j^{-1} requires only a small number of fast Fourier transforms. Furthermore, since M_j^{-1} only serves as a preconditioner, we do not need to be very accurate when applying M_j^{-1} . This allows us to further reduce the number of terms in the expansion of the symbol of M_j^{-1} .

Example 3. The sound speed $c(x)$ of this example is given in Figure 6. We perform the test on different combination of ω and N with ω/N fixed at about 16 points per wavelength. We compute the solutions using the BICGSTAB algorithm with relative error equal to 10^{-3} . The numerical results are summarized in Table 5. For each test, we report the number of iterations and the running time, both for the unconditioned system and the preconditioned system with M_1 and M_2 .




Figure 6: Sound speed $c(x)$ of Example 3.

Example 4. In this example, the sound speed $c(x)$ (shown in Figure 7) is a Gaussian bump. We perform the similar tests and the numerical results are summarized in Table 6.

In these two examples, we are able to use only 2 to 3 terms in the symbol expansion (15) of M_1^{-1} and M_2^{-1} . The results show that the preconditioners M_1 and M_2 reduce the number of iterations by a factor of 20 to 50, and the running time by a factor of 10 to 25. We also notice that the complex

$(w/2\pi, N)$	Uncond.		M_1		M_2	
	Iter	Time(s)	Iter	Time(s)	Iter	Time(s)
(4,64)	2243	8.40e+00	85	6.40e-01	57	5.10e-01
(8,128)	5182	6.79e+01	150	4.16e+00	88	2.46e+00
(16,256)	10412	6.50e+02	498	6.79e+01	354	4.82e+01
(32,512)			900	6.41e+02	306	2.20e+02

Table 5: Results of Example 3. For each test, we report the number of iterations and the running time in seconds.




Figure 7: Sound speed $c(x)$ of Example 4.

$(w/2\pi, N)$	Uncond.		M_1		M_2	
	Iter	Time(s)	Iter	Time(s)	Iter	Time(s)
(4,64)	3460	1.30e+01	67	5.00e-01	42	3.20e-01
(8,128)	10609	1.39e+02	210	5.80e+00	116	3.19e+00
(16,256)	35114	1.93e+03	1560	2.16e+02	681	9.56e+01
(32,512)			1550	1.12e+03	646	4.63e+02

Table 6: Results of Example 4. For each test, we report the number of iterations and the running time in seconds.

preconditioner M_2 outperforms M_1 by a factor of 2. This is in line with observations made in [16], where the complex constant appearing in front of the $\omega^2/c^2(x)$ term in M_1 and M_2 was optimized.

In these examples, we have not made the effort to optimize the coefficients in the preconditioners M_1 and M_2 and the BICGSTAB algorithm might not be the best iterative solver for this type of the problem. Here, our goal in these examples is to demonstrate that the DSC approach can offer an alternative to the multigrid method for the shifted Laplace preconditioner. Let us also note that we only consider the complex-shifted Laplace preconditioner in isolation, without implementing any additional deflation technique. Those seem to be very important in practice [17].

4.3 Polarization of wave operator

Another application of the discrete symbol calculus is to “polarize” the initial condition of linear hyperbolic systems. We consider the second order wave equation with variable coefficients

$$\begin{cases} u_{tt} - \operatorname{div}(\alpha(x)\nabla u) = 0 \\ u(0, x) = u_0(x) \\ u_t(0, x) = u_1(x) \end{cases}$$

with the extra condition $\int u_1(x)dx = 0$. Since the operator $L := -\operatorname{div}(\alpha(x)\nabla)$ is symmetric positive definite, its square root $P := L^{1/2}$ is well defined. We can use P to factorize the equation into

$$(\partial_t + iP)(\partial_t - iP)u = 0.$$

The solution $u(t, x)$ can be represented as

$$u(t, x) = e^{itP}u_+(x) + e^{-itP}u_-(x)$$

where the polarized components $u_+(x)$ and $u_-(x)$ of the initial condition are given by

$$u_+ = \frac{u_0 + (iP)^{-1}u_1}{2} \quad \text{and} \quad u_- = \frac{u_0 - (iP)^{-1}u_1}{2}.$$

We first use the discrete symbol calculus to compute the operator P^{-1} . Once P^{-1} is available, the computation of u_+ and u_- is straightforward.

Example 5. The coefficient $\alpha(x)$ in this example is shown in Figure 8 (a). The initial condition is set to be a plane wave solution of the unit sound speed:

$$u_0(x) = e^{2\pi ikx} \quad \text{and} \quad u_1(x) = -2\pi i|k|e^{2\pi ikx},$$

where k is a fixed wave number. If $\alpha(x)$ were equal to 1 everywhere, this initial condition itself would be polarized and the component $u_+(x)$ would be zero. However, due to the inhomogeneity in $\alpha(x)$, we expect both u_+ and u_- to be non-trivial after the polarization. The real part of $u_+(x)$ is plotted in Figure 8 (b). We notice that the amplitude $u_+(x)$ scales with the difference between the coefficient $\alpha(x)$ and 1. This is compatible with the asymptotic analysis of the operator P for large wave number. The figure of $u_-(x)$ is omitted as visually it is close to $u_0(x)$.

Example 6. The coefficient $\alpha(x)$ here is a random bandlimited function shown in Figure 9 (a). The initial conditions are the same as the ones used in Example 5. The real part of the polarized component $u_+(x)$ is shown in Figures 9 (b). Again, we see that the dependence of the amplitude of $u_+(x)$ on the difference between $\alpha(x)$ and 1.

4.4 Seismic depth migration

The setup is the same as in the Introduction: consider the Helmholtz equation

$$u_{zz} + \Delta_{\perp} + \frac{\omega^2}{c^2(x)}u = 0 \tag{16}$$

for $z \geq 0$. The transverse variables are either $x \in [0, 1]$ in the 1D case, or $x \in [0, 1]^2$ in the 2D case. (Our notations support both cases.) Given the wave field $u(x, 0)$ at $z = 0$, we want to compute the




Figure 8: Example 5. The real part of the polarized component $u_+ = (u_0 + (iP)^{-1}u_1)/2$. Notice that the amplitude of $u_+(x)$ scales with the quantity $\alpha(x) - 1$. $u_- = (u_0 - (iP)^{-1}u_1)/2$ is omitted since visually it is close to u_0 .




Figure 9: Example 6. The real part of the polarized component $u_+ = (u_0 + (iP)^{-1}u_1)/2$. Notice that the amplitude of $u_+(x)$ scales with the quantity $\alpha(x) - 1$. $u_- = (u_0 - (iP)^{-1}u_1)/2$ is omitted since visually it is close to u_0 .

wavefield for $z > 0$. For simplicity, we consider periodic boundary conditions in x or (x, y) , and no right-hand side in (16).

As mentioned earlier, we wish to solve the corresponding SSR equation

$$\left(\frac{\partial}{\partial z} - B(z) \right) u = 0, \quad (17)$$

where $B(z)$ is a regularized square root of $-\Delta_{\perp} - \omega^2/c^2(x, z)$. Call ξ the variable(s) dual to x . The locus where the symbol $4\pi^2|\xi|^2 - \omega^2/c^2(x, z)$ is zero is called the characteristic set of that

symbol; it poses well-known difficulties for taking the square root. To make the symbol elliptic (here, negative) we simply introduce

$$a(z; x, \xi) = g \left(4\pi^2 |\xi|^2, \frac{1}{2} \frac{\omega^2}{c^2(x, z)} \right) - \frac{\omega^2}{c^2(x, z)},$$

where $g(x, M)$ is a smooth version of the function $\min(x, M)$. Call $b(z; x, \xi)$ the symbol-square-root of $a(z; x, \xi)$, and $\tilde{B}(z) = b(z; x, i\nabla_x)$ the resulting operator. A large-frequency cutoff now needs to be taken to correct for the errors introduced in modifying the symbol as above. Consider a function $\chi(x)$ equal to 1 in $(-\infty, -2]$, and that tapers off in a C^∞ fashion to zero inside $[-1, \infty)$. We can now consider $\chi(b(z; x, \xi))$ as the symbol of a smooth “directional” cutoff, defining an operator $X = \chi(b(z; x, -i\nabla_x))$ in the standard manner. The operator $\tilde{B}(z)$ should then be modified as

$$X\tilde{B}(z)X.$$

At the level of symbols, this is of course $(\chi(b)) \# b \# (\chi(b))$ and should be realized using the composition routine of discrete symbol calculus.

Once this modified square root has been obtained, it can be used to solve the SSR equation. Formally, the operator mapping $u(x, 0)$ to $u(x, z)$ can be written as

$$E(z) = \left(\exp \int_0^z B(s) ds \right)_+,$$

where $(\exp \cdot)_+$ denotes the ordered exponential. If $B(s)$ were to make sense, this formula would be exact. Instead, we substitute $X\tilde{B}(s)X$ for $B(s)$, use the usual operator exponential as a simplification, and compute $E(z)$ using discrete symbol calculus. We intend for z to be small, i.e., comparable to the wavelength of the field $u(x, 0)$, in order to satisfy a CFL-type condition. With this type of restriction on z , the symbol of $E(z)$ remains sufficiently smooth for the DSC algorithm to be efficient⁷: the integral over s can be discretized by a quadrature over a few points, and the operator exponential is a good approximation to the ordered exponential that can be realized by scaling-and-squaring.

The effect of the cutoffs X is to smoothly remove 1) turning rays, i.e., waves that would tend to travel in the horizontal direction or even overturn, and 2) evanescent waves, i.e., waves that decay exponentially in z away from $z = 0$. This is why X is called a directional cutoff. It is important to surround \tilde{B} with *two* cutoffs to prevent the operator exponential from introducing energy near the characteristic set of the generating symbol $4\pi^2|\xi|^2 - \omega^2/c^2(x, z)$. This precaution would be hard to realize without an accurate way of computing compositions (twisted product). Note that the problem of controlling the frequency leaking while taking an operator exponential was already addressed by Chris Stolk in [48], and that our approach provides another, clean solution.

We obtain the following numerical examples.

Example 7. Let us start by considering the 1D case. The sound speed $c(x)$ in this example is a Gaussian waveguide (see Figure 10 (a)). We set ω to be $100 \cdot 2\pi$ in this case.

We perform two tests in this example. In the first test, we select the boundary condition $u(x, 0)$ to be equal to one. This corresponds to the case of a plane wave entering the waveguide. The solution of (17) is shown in Figure 10 (b). As z grows, the wave front starts to deform and the caustics appears at $x = 1/2$ when the sound speed $c(x)$ is minimum.

⁷For larger $E(z)$ would be a Fourier integral operator, and a phase would be needed in addition to a symbol. We leave this to a future project.

In the second test of this example, we choose the boundary condition $u(x, 0)$ to be a Gaussian wave packet localized at $x = 1/2$. The wave packet enters the wave guide with an incident angle about 45 degrees. The solution is shown in Figure 10 (c). Even though the wave packet deforms its shape as it travels down the wave guide, it remains localized. Notice that the packet bounces back and forth at the regions with large sound speed $c(x)$, which is the result predicted by geometric optics in the high frequency regime.




Figure 10: Example 7. (a) sound speed $c(x)$. (b) the solution when the boundary condition $u(x, 0)$ is a constant. (c) the solution when the boundary condition $u(x, 0)$ is a wave packet.

Example 8. Let us now consider the 2D case. The sound speed used here is a two dimensional Gaussian waveguide (see Figure 11 (a)). We again perform two different tests. In the first test, the boundary condition $u(x, y, 0)$ is equal to a constant. The solution at the cross section $y = 1/2$ is shown in Figure 11 (b). In the second test, we choose the boundary condition to be a Gaussian wave packet with oscillation in the x direction. The packet enters the waveguide with an incident angle of 45 degrees. The solution at the cross section $y = 1/2$ is shown in Figure 11 (c). Both of these results are similar to the ones of the one dimensional case.




Figure 11: Example 8. (a) sound speed $c(x)$. (b) the solution at the cross section $y = 1/2$ when the boundary condition $u(x, y, 0)$ is a constant. (c) the solution at the cross section $y = 1/2$ when the boundary condition $u(x, y, 0)$ is a wave packet.

5 Discussion

5.1 Other domains and boundary conditions

An interesting question is what form discrete symbol calculus should take when other boundary conditions than periodic are considered, or on more general domains than a square.

One can speculate that the discrete Sine transform (DST) should be used as e_λ for Dirichlet boundary conditions on a rectangle, or the discrete Cosine transform (DCT) for Neumann on a rectangle. Whatever choice is made for e_λ should dictate the definition of the corresponding frequency variable ξ . A more robust approach could be to use spectral elements for more complicated domains, where the spectral domain would be defined by Chebyshev expansions. One may also imagine expansions in prolate spheroidal wavefunctions. Regardless of the type of expansions chosen, the theory of pseudodifferential operators on bounded domains is a difficult topic that will need to be understood.

For applications to wave propagation, it will be important to generate symbols that handle non-reflecting boundary conditions, perhaps implemented as a PML [2, 33]. Such boundary conditions make the spectrum complex, hence would substantially complicate the square root operation. In fact, we hope that ideas of symbol calculus will themselves provide new insights into solving the problem of accurate nonreflecting boundary conditions in variable media.

5.2 Other equations

Symbol-based methods may help solve other equations than elliptic PDE. The heat equation in variable is one of them: its fundamental solution has a nice pseudodifferential smoothing form that can be computed via scaling-and-squaring.

A more challenging example are hyperbolic systems in variable, smooth media. The time-dependent Green's function of such systems is not a pseudodifferential operator, but rather a Fourier integral operator (FIO), where $e^{2\pi i x \cdot \xi} a(x, \xi)$ needs to be replaced by $e^{\Phi(x, \xi)} a(x, \xi)$. We regard the extension of discrete symbol calculus to handle such phases a very interesting problem, see [11, 12] for preliminary results on fast application of FIO.

A Appendix

Proof of Lemma 1. Set

$$\hat{a}_\lambda(\xi) = \int e^{-2\pi i x \cdot \lambda} a(x, \xi) dx$$

for the Fourier series coefficients of $a(x, \xi)$ in x . Then we can express (6) as

$$(Af)(x) = \sum_{\xi \in \mathbb{Z}^d} e^{2\pi i x \cdot \xi} \sum_{\lambda \in \mathbb{Z}^d} e^{2\pi i x \cdot \lambda} \hat{a}_\lambda(\xi) \hat{f}(\xi).$$

We seek to interchange the two sums. Since $a(x, \xi)$ is differentiable d' times, we have

$$(1 + |2\pi\lambda|^{d'}) \hat{a}_\lambda(\xi) = \int_{[0,1]^d} e^{-2\pi i x \cdot \lambda} (1 + (-\Delta_x)^{d'/2}) a(x, \xi) dx,$$

hence $|\hat{a}_\lambda(\xi)| \leq (1 + |2\pi\lambda|^{d'})^{-1} \|(1 + (-\Delta_x)^{d'/2}) a(x, \xi)\|_{L_x^\infty}$. The exponent d' is chosen so that $\hat{a}_\lambda(\xi)$ is absolutely summable in $\lambda \in \mathbb{Z}^d$. If in addition we assume $\hat{f} \in \ell_1(\mathbb{Z}^d)$, then we can apply Fubini's theorem and write

$$(Af)(x) = \sum_{\lambda \in \mathbb{Z}^d} A^\lambda f(x),$$

where $A^\lambda f(x) = e^{2\pi i x \cdot \lambda} (M_{\hat{a}_\lambda(\xi)} f)(x)$, and M_g is the operator of multiplication by g on the ξ side. By Plancherel, we have

$$\|A^\lambda f\|_{L^2} = \|M_{\hat{a}_\lambda(\xi)} f\|_{L^2} \leq \sup_{\xi} |\hat{a}_\lambda(\xi)| \cdot \|f\|_{L^2}.$$

Therefore, by the triangle inequality,

$$\begin{aligned} \|Af\|_{L^2} &\leq \sum_{\lambda \in \mathbb{Z}^d} \|A^\lambda f\|_{L^2} \\ &\leq \sum_{\lambda \in \mathbb{Z}^d} (1 + |2\pi\lambda|^{d'})^{-1} \cdot \sup_{x, \xi} |(1 + (-\Delta_x)^{d'/2}) a(x, \xi)| \cdot \|f\|_{L^2}. \end{aligned}$$

As we have seen, the sum over λ converges. This proves the theorem when f is sufficiently smooth; a classical density argument shows that the same conclusion holds for all $f \in L^2([0, 1]^d)$. \square

References

- [1] G. Bao and W. Symes. Computation of pseudo-differential operators. *SIAM J. Sci. Comput.*, 17(2):416–429, 1996.
- [2] J. P. Berenger, A perfectly matched layer for the absorption of electromagnetic waves, *J. Comput. Phys.* 114(2):185–200, 1994.
- [3] G. Beylkin, R. Coifman, and V. Rokhlin. Fast wavelet transforms and numerical algorithms. I. *Comm. Pure Appl. Math.*, 44(2):141–183, 1991.
- [4] G. Beylkin and M. J. Mohlenkamp. Numerical operator calculus in high dimensions. *Proc. Nat. Acad. Sci.*, 99(16):10246–10251, 2002.
- [5] G. Beylkin and M. J. Mohlenkamp. Algorithms for numerical analysis in high dimensions. *SIAM J. Sci. Comput.*, 26(6):2133–2159, 2005.
- [6] G. Beylkin, M. J. Mohlenkamp, and F. Perez. Approximating a wavefunction as an unconstrained sum of Slater determinants. *J. Math. Phys.*, 49:032107, 2008.
- [7] G. Beylkin and K. Sandberg. Wave propagation using bases for bandlimited functions. *Wave Motion*, 41:263–291, 2005.
- [8] J. P. Boyd. *Chebyshev and Fourier spectral methods*. Dover Publications Inc., Mineola, NY, second edition, 2001.
- [9] S. Börm, L. Grasedyck, and W. Hackbusch. Hierarchical matrices. Technical Report 21, Max-Planck-Institut für Mathematik in den Naturwissenschaften, Leipzig, 2003.
- [10] E. J. Candès, L. Demanet, D. L. Donoho and L. Ying. Fast discrete curvelet transforms. *SIAM Multiscale Model. Simul.*, 5(3):861–899, 2006.
- [11] E. J. Candès, L. Demanet and L. Ying. Fast computation of Fourier integral operators. *SIAM J. Sci. Comput.*, 29(6):2464–2493, 2007.
- [12] E. J. Candès, L. Demanet and L. Ying. Optimal computation of Fourier integral operators via the Butterfly algorithm. Submitted, 2008.
- [13] L. Demanet. Curvelets, Wave Atoms, and Wave Equations. *Ph.D. Thesis*, California Institute of Technology, 2006.
- [14] L. Demanet and L. Ying. Wave atoms and time upscaling of wave equations. *Numer. Math.*, to appear, 2008.
- [15] Y. Erlangga. Advances in iterative methods and preconditioners for the Helmholtz equation. *Arch. Comput. Methods. Eng.*, 15:37–66, 2008.
- [16] Y. Erlangga. A robust and efficient iterative method for the numerical solution of the Helmholtz equation. *PhD Thesis*, Delft University, 2005.
- [17] Y. Erlangga and R. Nabben. Multilevel Projection-Based Nested Krylov Iteration for Boundary Value Problems. *SIAM J. Sci. Comput.*, 30(3):1572–1595, 2008.
- [18] L. Fishman, M. V. de Hoop, and M. van Stralen. Exact constructions of square-root Helmholtz operator symbols: The focusing quadratic profile. *J. Math. Phys.* 41(7):4881–4938, 2000.

- [19] M. D. Fisk and G. D. McCartor. The phase screen method for vector elastic waves. *J. Geophys. Research*, 96(B4):5985–6010, 1991.
- [20] G. B. Folland. *Harmonic analysis in phase-space*. Princeton university press, 1989.
- [21] M. Gander and F. Nataf. An incomplete LU preconditioner for problems in acoustics. *J. Comput. Acoust.* 13(3):455–476, 2005.
- [22] L. Greengard and V. Rokhlin. A fast algorithm for particle simulations. *J. Comput. Phys.*, 73:325, 1987.
- [23] S. Chandrasekaran and M. Gu. A fast and stable solver for recursively semi-separable systems of linear equations, in *Structured matrices in mathematics, computer science, and engineering*, Contemporary Mathematics 281, 2001.
- [24] W. Hackbusch. A sparse matrix arithmetic based on H -matrices. I. Introduction to H -matrices. *Computing*, 62:89–108, 1999.
- [25] W. Hackbusch and Z. P. Nowak. On the fast matrix multiplication in the boundary element method by panel clustering. *Numer. Math.*, 54:463–491, 1989.
- [26] F. J. Herrmann, P. P. Moghaddam and C. C. Stolk. Sparsity- and continuity-promoting seismic image recovery with curvelet frames. *Appl. Comput. Harmon. Anal.* 24(2):150–173, 2008.
- [27] N. J. Higham. *Functions of matrices: Theory and Computation*. SIAM, 2008.
- [28] N. J. Higham. Stable iterations for the matrix square root. *Numer. Algorithms*, 15:227–242, 1997.
- [29] L. Hörmander. *The Analysis of Linear Partial Differential Operators*. 4 volumes, Springer, 1985.
- [30] A. S. Householder. *The Theory of Matrices in Numerical Analysis*. Dover Publications, 2006.
- [31] T. Huckle. Compact Fourier Analysis for Designing Multigrid Methods *SIAM J. Sci. Comput.* 31(1):644–666, 2008.
- [32] 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):35–51, 1994.
- [33] D. Komatitsch and J. Tromp. A perfectly matched layer absorbing boundary condition for the second-order seismic wave equation *Geophys. J. Int.* 154(1):146–153, 2003.
- [34] M. P. Lamoureux and G. F. Margrave. An Introduction to Numerical Methods of Pseudodifferential Operators. *Proc. CIME Workshop on Pseudodifferential Operators, Quantization and Signals*, 2006
- [35] J. H. Le Rousseau and M. V. de Hoop. Generalized-screen approximation and algorithm for the scattering of elastic waves. *Q. J. Mech. Appl. Math.* 56:1–33, 2003.
- [36] T. Lin and F. Herrmann. Compressed wavefield extrapolation. *Geophysics*, 72(5):77–93, 2007.
- [37] Y. Meyer. Wavelets and operators. Analysis at Urbana, *London Math. Soc. Lecture Notes Series*, 137:256–364, Cambridge Univ. Press, 1989.

- [38] E. Michielssen and A. Boag. A multilevel matrix decomposition algorithm for analyzing scattering from large structures. *IEEE Trans. Antennas Propag.* 44(8):1086–1093, 1996.
- [39] C. Moler and C. Van Loan. Nineteen Dubious Ways to Compute the Exponential of a Matrix, Twenty-Five Years Later. *SIAM Review* 45(1):3–49, 2003.
- [40] G. Peyré and S. Mallat. Orthogonal Bandlet Bases for Geometric Images Approximation. *Comm. Pure Appl. Math.*, to appear, 2008.
- [41] G. Schulz. Iterative Berechnung der reziproken Matrix. *Z. Angew. Math. Mech.* 13:57–59, 1933.
- [42] R. T. Seeley. Complex powers of an elliptic operator. *Proc. Symp. Pure Math.*, 10:288–307, 1967.
- [43] S. Serra. Optimal, quasi-optimal and superlinear band-Toeplitz preconditioners for asymptotically ill-conditioned positive definite Toeplitz systems. *Math. Comp.* 66(218):651–665, 1997.
- [44] M. A. Shubin. Almost periodic functions and partial differential operators. *Russian Math. Surveys* 33(2):1–52, 1978.
- [45] C. Sogge. *Fourier Integrals in Classical Analysis*. Cambridge University Press, 1993.
- [46] E. Stein. *Harmonic Analysis*. Princeton University Press, 1993.
- [47] C. C. Stolk. A fast method for linear waves based on geometrical optics. Preprint, 2007.
- [48] C. C. Stolk. A pseudodifferential equation with damping for one-way wave propagation in inhomogeneous acoustic media. *Wave Motion* 40(2):111–121, 2004.
- [49] M. Taylor. *Pseudodifferential Operators and Nonlinear PDE*. Birkäuser, Boston, 1991.
- [50] L. N. Trefethen. *Spectral methods in MATLAB*, volume 10 of *Software, Environments, and Tools*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2000.