Curvelets, Wave Atoms, and Wave Equations

Thesis by

Laurent Demanet

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

California Institute of Technology

Pasadena, California

2006

(Defended May 19, 2006)

Acknowledgements

I am deeply grateful to my advisor Emmanuel Candès for helping create the best possible academic environment I could have hoped for during my four years at Caltech. His scientific vision, enthusiasm and brilliance have been nothing short of a model for me. I like to believe that Emmanuel's deep appreciation of mathematics has played a role in defining my "taste" and, I hope, will inspire many others as well. The research directions that we have explored together are in my own (biased) opinion very exciting; and for this I feel very obliged. I also wish to thank him for his kind and sometimes humbling encouragements.

The next person who has influenced my scientific life in no small part, and a great friend, is Lexing Ying. I thank him wholeheartedly for sharing his wide expertise with me, for the passionate discussions we engaged in together, and for his infinite patience with the computer implementations that he kindly allowed me, and helped me to report on in this thesis.

I would like to express my gratitude to Thomas Hou, Jerrold Marsden, and Houman Owhadi for serving on my thesis committee, as well as all the professors whom I took classes from at Caltech, for their uniform excellence in teaching. Special thoughts go to Oscar Bruno, Wilhelm Schlag, and Jean-Pierre Antoine, whose teaching skills I greatly admire.

During the fall of 2004 I took part in the Multiscale Geometric Analysis program at IPAM, UCLA, and I wish to thank all the participants and organizers for making the events so enjoyable. I was very fortunate to meet or further interact with Felix Herrmann, William Symes, Maarten de Hoop, and many others while at IPAM.

The scientific atmosphere in Caltech's Applied and Computational Mathematics department is terrific. Thank you to the student chapter of SIAM, and in particular my friends Jon, Stéphane, Hannes, Marco and Paige for their dedication in organizing our famous events. Thank you to the rest of the Candès group, and in particular Justin Romberg, for

the invaluable group meeting discussions. Thank you to my officemates Wuan and Marco for bearing with me for the past 3 years. Finally, a million thanks to Sheila, Chad and Sydney for their remarkable work and for helping make the department what it is.

Although an ocean and then some separates me from my parents, family and friends in Belgium and elsewhere in Europe, please be assured that you are very close in my heart! I am also grateful to all my local friends for making my stay in Pasadena all the more pleasant, in particular Kai, Sébastien, Tonci, Grace and Lars.

I have been truly blessed to meet Young two and a half years ago and one sentence cannot do justice to the way she has forever changed my life.

Abstract

We argue that two specific wave packet families—curvelets and wave atoms—provide powerful tools for representing linear systems of hyperbolic differential equations with smooth, time-independent coefficients. In both cases, we prove that the matrix representation of the Green's function is

- *sparse* in the sense that the matrix entries decay nearly exponentially fast (i.e., faster than any negative polynomial), and
- *well organized* in the sense that the very few nonnegligible entries occur near a few shifted diagonals, whose location is predicted by geometrical optics.

This result holds only when the basis elements obey a precise parabolic balance between oscillations and support size, shared by curvelets and wave atoms but not wavelets, Gabor atoms, or any other such transform.

A physical interpretation of this result is that curvelets may be viewed as coherent waveforms with enough frequency localization so that they behave like waves but at the same time, with enough spatial localization so that they simultaneously behave like particles.

We also provide fast digital implementations of tight frames of curvelets and wave atoms in two dimensions. In both cases the complexity is $O(N^2 \log N)$ flops for N-by-N Cartesian arrays, for forward as well as inverse transforms.

Finally, we present a geometric strategy based on wave atoms for the numerical solution of wave equations in smoothly varying, 2D time-independent periodic media. Our algorithm is based on sparsity of the matrix representation of Green's function, as above, and also exploits its *low-rank block structure* after separation of the spatial indices. As a result, it becomes realistic to accurately build the full matrix exponential using repeated squaring, up to some time which is much larger than the CFL timestep. Once available, the wave atom representation of the Green's function can be used to perform 'upscaled' timestepping.

We show numerical examples and prove complexity results based on a priori estimates of sparsity and separation ranks. They beat the $O(N^3)$ bottleneck on an N-by-N grid, for a wide range of physically relevant situations. In practice, the current wave atom solver can become competitive over a pseudospectral method in the regime when the wave equation should be solved several times with different initial conditions, as in reflection seismology.

Contents

Chapter 1

Introduction

The main object of study in this thesis is the wave equation in two or three space dimensions,

$$\frac{\partial^2 u}{\partial t^2}(t,x) = c^2(x)\Delta u(t,x), \qquad u(0,x) = u_0(x), \qquad \frac{\partial u}{\partial t}(0,x) = u_1(x), \tag{1.1}$$

a good model for acoustic and electromagnetic waves. For simplicity, we consider a spatial domain without boundaries. When $c(x)$ is smooth, solutions to the wave equation are waves traveling with local speed $c(x)$.

The classical mathematical questions of existence, uniqueness, and Strichartz-type estimates were settled long ago for this simple equation [38], yet the basic observation that wave equations model traveling waves is surprisingly absent from their justifications. Standard proof techniques involve energy estimates and weak convergence arguments. In fact, Lions and Magenes [60], and also Stolk [80] use these analysis tools to prove existence and uniqueness in the general case $\log c(x) \in L^\infty$, a regime of rough coefficients in which solutions may not behave like traveling waves at all.

Likewise, textbook numerical methods for the wave equation in two or three space dimensions tend to ignore the geometry of wave propagation. Degrees of freedom are usually distributed in a mechanical way over some grid points (for finite differences), grid cells (for finite volumes), spectral or other Galerkin elements, leaving for the physics of the problem to be discovered *a posteriori*, from the output of a computer simulation. In this spirit, solving say the heat equation instead would require little modification of the source code.

The message of the present thesis is that the geometric aspects of wave propagation are very well understood mathematically and make for excellent *a priori* geometric information for novel numerical methods. Computational wave propagation in smooth media is, for

example, of prime importance in seismic imaging.

We devote the next section to retracing some key steps in the genesis of an exciting dialogue between pure and applied mathematics. Three surprisingly interconnected problems form the mathematical context of this thesis: propagation of singularities, boundedness estimates, and multiscale compression of operators. We make no claim of reviewing the literature in any exhaustive manner. Sections 1.2 through 1.4 summarize the contributions of this thesis and introduce Chapters 2 through 4 respectively.

1.1 Context

1.1.1 Propagation of Singularities

There is a slogan in textbooks on partial differential equations (PDE), that solutions to hyperbolic equations have singularities that propagate along characteristics [38]. An accessible formal justification of this result can be found in [89]. The mathematical formulation of the problem and its resolution for linear equations are part of a much larger body of theory developed mostly in the 1970s and 1980s, called microlocal analysis.

To see how microlocal analysis gives information about the wave equation (1.1), let us express its solution $u(t)$ by means of integral operators with kernels G_0 and G_1, also called retarded propagator or Green's function,

$$u(t,x) = \int_{\mathbb{R}^n} G_0(t,x,y)u_0(y)\,dy + \int_{\mathbb{R}^n} G_1(t,x,y)u_1(y)\,dy.$$

If (1.1) were replaced by some other well-posed linear initial-value PDE, the singularities of u_0 and u_1 would most likely be washed out by diffusion or dispersion. But because we are in presence of the wave equation, singularities propagate in a very predictable manner. This property reflects itself in the fact that the two kernels G_0 and G_1 are smooth almost everywhere, but share the same locus of singularity, a codimension one manifold $x \in \Gamma_y(t)$ for each time $t > 0$ and $y \in \mathbb{R}^n$.

In 1957 Peter Lax showed in a pioneering article [57] that G_0 and G_1 can be well approximated for small times by oscillatory integrals, now called Fourier integral operators, as

$$G_j(t,x,y) \simeq \sum_{k=\pm} \int_{\mathbb{R}^n} e^{i(-y\cdot\xi+\Phi_k(t,x,\xi))} a_{j,k}(x,\xi)\,d\xi, \qquad j = 0, 1. \tag{1.2}$$

The near-equality sign is to be understood in the sense that the difference between each G_j and its integral approximation is smoother than either of them. The phases $\Phi_\pm$ solve the Hamilton-Jacobi equations of geometrical optics,

$$\frac{\partial \Phi_\pm}{\partial t}(t,x,\xi) = \pm c(x)|\nabla_x \Phi_\pm(t,x,\xi)|, \qquad \Phi_\pm(0,x,\xi) = x\cdot\xi, \tag{1.3}$$

which are also the characteristic equations for (1.1). The amplitudes a_{jk} belong to some standard symbol class that will concern us later. The integral in (1.2) is then expected to diverge when the phase in the integrand's exponent is stationary, providing a characterization of the singularities of G_0, G_1 as

$$\Lambda_y(t) = \{x \in \mathbb{R}^n : \text{ there exists } \xi \in \mathbb{R}^n \text{ and a choice of sign for which } y = \nabla_\xi \Phi_\pm(t,x,\xi)\}.$$

In the case when $c(x)$ is constant, for example, $\Phi_\pm(t,x,\xi) = x\cdot\xi \pm ct|\xi|$ and $\Lambda_y(t)$ is the boundary of the forward light cone with apex at y.

Lax's original approach is compelling but was soon recast by Hörmander, Duistermaat, and others [12] into a more powerful, far-reaching 'phase-space' viewpoint, later dubbed microlocal analysis. The rationale of their approach is that *directions* in which a function or distribution is singular matter as much as the *location* of the singularity. Mathematically, directional smoothness is encoded in the notion of *wavefront set* of a distribution, defined as the complement of the set of directions in which the Fourier transform of the distribution, properly localized near x, decays fast. For example, if we consider the Green's function of the wave equation in frequency,

$$\hat{G}_j(t,\eta,\xi) \simeq \sum_{k=\pm} \int_{\mathbb{R}^n} e^{i(-x\cdot\eta + \Phi_k(t,x,\xi))} a_{j,k}(x,\xi)\, dx, \qquad j = 0,1,$$

then this time, stationary phase analysis yields the relation $\eta = \nabla_x \Phi_\pm(t,x,\xi)$. In case G_j is singular at (x,y), we can gather additional information on the *directions* (η,ξ) in which $G_j(x,y)$ fails to be smooth, yielding the wavefront set of $G_j(x,y)$,

$$\begin{aligned}\Lambda(t) = \{(x,\eta) \in \mathbb{R}^{2n} : &\text{ there exists } (y,\xi) \in \mathbb{R}^{2n} \text{ and a choice of sign for which} \\ &y = \nabla_\xi \Phi_k(t,x,\xi) \text{ and } \eta = \nabla_x \Phi_k(t,x,\xi)\}.\end{aligned} \tag{1.4}$$

The above set is equivalently considered as the corresponding union of quadruples $(x, \eta, y, \xi) \in \mathbb{R}^{4n}$ and then called canonical relation.[1] The set $\mathbb{R}^{2n}$ of couples (x, η) is called cotangent bundle in mathematics, or phase-space in physics. The fundamental theorem of propagation of singularities is the following result [51].

A singularity at x, in the direction η gets mapped by the wave group to at most two points $y_\pm$ and directions $\xi_\pm$ specified by the canonical relation.

Note that consideration of the light cone $\Lambda_y(t)$ alone does not suffice to formulate such a precise result; the phase-space viewpoint was essential.

The dual variable ξ or η is also central in the Hamiltonian formulation of geometrical optics as a system of ordinary differential equations (ODE) for the "light" rays,

$$\begin{cases} \dot{x}(t) &= c(x(t))\frac{\xi(t)}{|\xi(t)|}, & x(0) &= x_0, \\ \dot{\xi}(t) &= -\nabla c(x)|\xi(t)|, & \xi(0) &= \xi_0. \end{cases} \tag{1.5}$$

These are the characteristic equations for the Hamilton-Jacobi equation (1.3), the trajectories of which are called "bicharacteristic strips." The union of the two prescriptions $(x_0, \xi_0) \to (x(t), \xi(t))$ for $k = \pm$, is for small times equivalent to the canonical relation (1.4) for propagation of singularities.

Remarkably the system (1.5) is solvable for all times, unlike the Hamilton-Jacobi equation (1.3) that generally breaks down upon formation of so-called caustics—because some $\Phi_\pm$ would become multivalued and could not be the solution of a PDE anymore. The canonical relation (1.4) defined for large times by means of (1.5) is a more intrinsic object than the phase functions themselves.[2]

Generalization of results of propagation of singularities to large times by means of a phase-space approach is the far-reaching consequence of microlocal analysis we were mentioning at the beginning of this section.

[1] For specialists, we are actually in presence of two canonical relations, each correponsing to a choice of sign. Also, canonical relations sometimes come with ξ replaced by $-\xi$ depending on the choice of symplectic form in the direct product of cotangent bundles, to make sure the canonical relation is a Lagrangian manifold. The term 'canonical' refers to this latter property. See [51].

[2] Artificial large-time complications occur only because the projection of $\Lambda(t)$ onto physical space may have folds, hence the presence of caustics as creases.

Let us however remark that most of the ideas behind geometrical optics and geometrical mechanics had been around for a long time before Hörmander wrote his first treatise, even beyond caustics. Phase-space has been a central notion in Hamiltonian mechanics for almost two hundred years. It is common knowledge since the 1930s that the Hamilton-Jacobi equation is naturally derived from a high-frequency ansatz for the wave equation, cf. WKB expansions in quantum mechanics. Some of the landmark papers on the high-frequency asymptotic analysis of wave equations in the neighborhood of caustics were produced in the 1950s, see Kay and Keller [54], and Ludwig [61].

In the past two decades microlocal analysis has emerged as a natural mathematical language in seismic imaging. Considerations of wave fronts and canonical relations are central in formulating and solving the inverse problem of determining discontinuities in the sound speed $c(x)$ (and other parameters) in the Earth's upper crust from acoustic wave measurements [6, 81, 82].

Microlocal ideas have also made a strong impact on contemporary mathematics, in particular on questions of propagation of singularities for nonlinear wave equations [10, 63, 30, 5], for wave equations on manifolds with edges or corners [58, 87], and many other problems.

1.1.2 Regularity of Integral Operators

In the late 1970s Antonio Córdoba and Charles Fefferman showed that the action of a Fourier Integral Operator (FIO) is simply a "rearrangement" of energy in phase-space, modulo negligible contributions [28]. Their study effectively casts the high-frequency viewpoint of singularity propagation into the bounds of the Heisenberg uncertainty principle, by means of a continuous wave packet representation. Important mathematical results follow from their analysis, in particular an elementary proof of L^2 boundedness of FIOs, by means of Schur's lemma.[3]

Schur's lemma is an instance of almost orthogonality method for proving boundedness, or regularity estimates, of operators. The Cotlar-Stein lemma is another such tool. Underlying both approaches is the key idea that, in order to understand the action of an operator and how it may or may not concentrate energy, it suffices to break it up into "almost orthogonal"

[3] A special case of Schur's lemma is the following: if one can find a domain indexed by λ, λ' in which the kernel T obeys $\sup_\lambda \sum_{\lambda'} |T(\lambda, \lambda')| < \infty$ and $\sup_{\lambda'} \sum_\lambda |T(\lambda, \lambda')| < \infty$, then T is bounded on L^2.

contributions, roughly independent of each other. For definitions and statements, see [79], Chapter 7.

By the early 1970s, Elias Stein and his collaborators had turned the boundedness question for singular integral (pseudodifferential) operators (ΨDO) into a branch of harmonic analysis [79], also called the Calderón program [65]. A significant milestone in this program is the T(1) theorem of David and Journé in 1984 [29], which gives necessary and sufficient conditions for L^2 boundedness in the so-called Calderón-Zygmund operator class. The correct argument in [29] was to define almost orthogonal blocks corresponding to different *scales*. In fact, one proof of their result is based on expressing almost orthogonality as diagonal dominance of the wavelet matrix

$$T(\lambda, \lambda') = \langle \psi_\lambda, T\psi_{\lambda'} \rangle, \tag{1.6}$$

where ψ_λ are wavelets, and λ indexes scales and positions. Some of the first constructions of orthonormal bases of wavelets by Y. Meyer [64] were actually motivated by the Calderón program.

With Fourier Integral Operators, boundedness questions are raised to a more difficult level where the geometry of wave propagation needs to be taken into account. Although L^2 boundedness is accessible by soft harmonic analysis arguments [79], it was only in 1991 that sharp L^p boundedness, $p \neq 2$, was established by Andreas Seeger, Chris Sogge, and Elias Stein [68] by means of interpolation from a Hardy-H^1 to L^1 endpoint estimate. In their argument they exhibit the correct decomposition of the Fourier domain, Fefferman's *second dyadic decomposition* (SDD), where dyadic annuli are further subdivided into thin rectangles obeying the parabolic scaling length = (width)^2. The SDD was first introduced in [39] to study boundedness of a family of Bochner-Riesz summation multipliers on L^4. Córdoba and Fefferman's decomposition is different from the SDD, yet shows some subtle resemblance.

While a standard dyadic decomposition of the frequency space corresponds to wavelet analysis, we now know that the second dyadic decomposition corresponds to *curvelet analysis* [15]. Curvelet were introduced as a tight frame of multiscale directional basis elements by Emmanuel Candès and David Donoho [18, 20], in an attempt to overcome inherent limitations of traditional multiscale representations such as wavelets. They efficiently address

important problems where wavelet ideas are far from ideal, like sparse representation of images with edges. Curvelets have now become a household tool in applied communities, for instance in geometrical image processing [76], seismic imaging [49], and computerized tomography.

A wave packet frame identical to curvelets was independently introduced by Hart Smith as a tool to define function spaces adapted to FIOs [73]. Smith also used curvelets to define a parametrix and formulate Strichartz and Pecher estimates for wave equations in rough ($C^{1,1}$) metrics [74]. More recently, Terence Tao applied a variation of the SDD to proving a weak-L^1 bound for FIOs [83].

The definitions of wavelets, curvelets, and the related decompositions of the Fourier domain will be given in Section 1.2.4 and Chapter 2.

1.1.3 Multiscale Compression of Operators

The beauty of almost orthogonality methods in harmonic analysis is their application as *compression* tools for operators, by revealing their true information content. It was recognized in 1991 by Gregory Beylkin, Ronald Coifman, and Vladimir Rokhlin [7] that Meyer's almost diagonal wavelet representation of singular integral operators gives rise to matrices that are well suited for numerical computations. For most pseudodifferential operators it suffices to put to zero the small matrix entries, below some threshold ϵ in absolute value, to obtain a sparse matrix with at most $O(\epsilon^{-1/M})$ elements per row and column for arbitrary large M—an optimal compression estimate.[4]

The wavelet sparsity result has been used extensively in the 1990s as the basis of innovative numerical methods for linear elliptic PDE and some boundary integral equations. See for instance [53, 27] and the book [26]. Beyond sparsity, wavelets also allow multiscale preconditioning of differential operators by simple diagonal matrices; an important asset for iterative inversion algorithms like conjugate gradients. As a result of these good properties wavelet-based numerical methods often enjoy asymptotically optimal complexity estimates. One could speak of wavelets as a universally parsimonious way of distributing degrees of freedom in the discretization of smooth elliptic problems. Wavelets have also been successfully applied to homogenization [34], solving the variable coefficient heat equation [37], 1D

[4]This claim is a theorem provided the amplitude of the ΨDO is in a symbol class $S^0_{\rho,\delta}$ with $\delta \le \rho \le 1$ and $\delta < 1$.

wave equations [37, 4], and many other problems.

In contrast, and for good reasons, oscillatory integrals and FIOs in dimensions two and greater do not lend themselves to a sparse wavelet representation.

For instance, the Green's function of wave equations yields poorly structured wavelet matrices because of geometric dispersion: wavelets do not remain wavelet-like waveforms but disperse in all directions as time increases. An illustration of this problem can be found in the introduction of our paper [15]. Accordingly, most efforts in the applied literature (although very interesting) go into designing adequate quadratures for oscillatory integrands without tapping into multiscale ideas [52]. Notable exceptions are [2]; and also [11, 25] in a different context, the Helmholtz equation.

To resolve this stalemate, we need to rethink the geometry of multiscale representations.

We show that microlocal and harmonic analysis, cf. the previous sections, provide the correct insights towards understanding compressibility of wave-type oscillatory integrals. Chapter 2 of the present thesis establishes this claim in a quantitative manner. There are essentially only two families of wave packets, or tilings of phase space, which provide a change of basis to make wave propagators asymptotically optimally sparse. These wave packet families, *curvelets* and *wave atoms*, inherit respectively the geometrical properties of the second dyadic decomposition and Córdoba-Fefferman's wave packets. In the spirit of Hörmander's large time theory, our phase-space constructions are completely oblivious to the presence of caustics.

The compression gains achieved in the right representation offer the tantalizing perspective that solving the wave equation itself could be thought of as an operator compression problem. The road from a sparsity theorem to a fast algorithm in the form of competitive software is nevertheless long and challenging. The requirements of an efficient harmonic-analysis-based method for operator compression are typically at least threefold, and include

- the availability of a fast digital transform for analysis into, and synthesis from coefficients; and
- the ability to predict the *location* and compute the *value* of significant matrix entries, in an efficient manner (one instance of such operation in the context of classical wavelet solvers is called the refinement rule); and
- a deliberate *pruning* strategy to discard the less important information. This step can

take the form of thresholding small entries in a matrix, resulting in a provably good approximation.

We will come back to these three important steps in Sections 1.3 and 1.4 and summarize the solutions we propose for each of them. As we will see, they do not automatically follow from conventional wavelet wisdom. The detailed exposition follows in Chapters 3 and 4.

To conclude this section let us remark that the field of scientific computing is alive and well. A few geometrically informed algorithms have been proposed in the recent years for linear wave and related equations. This list is, of course, far from complete.

- The work by de Hoop, le Rousseau, and Wu on phase-screen methods addresses fast computation of one-way (down-going or up-going) wave equations by means of adequate approximation of pseudodifferential symbols, with applications to seismic imaging [50].

- Beylkin and Sandberg [9] achieve remarkable accuracy in solving the time-dependent wave equation (1.1) using a basis of prolate spheroidal wavefunctions, well suited for approximation of bandlimited functions in intervals and encoding of boundary conditions on rectangular domains.

- Impressive progress has been made on the front of fast multipole methods for the Helmhotz equation in free space, by Rokhlin and co-workers. See [25].

- Achi Brandt presents the state of the art on multiscale ideas for a variety of equations in [11].

1.2 Curvelets, Wave Atoms, and Sparse Representations

In this section we will consider slightly more general initial value problems than (1.1), namely hyperbolic systems of the form

$$\frac{\partial u}{\partial t} + \sum_k A_k(x) \frac{\partial u}{\partial x_k} + B(x)u = 0, \qquad u(0,x) = u_0(x), \tag{1.7}$$

where u is an m-dimensional vector and $x \in \mathbb{R}^n$. The matrices A_k and B may depend on the spatial variable x, and the A_k are symmetric. Second-order wave equations like (1.1)

can be reduced to a symmetric system of first-order equations (1.7) by a standard change of variables.

Linear hyperbolic systems are ubiquitous in the sciences and a classical example are the equations for acoustic waves, Maxwell's equations of electrodynamics and the equations of linear elasticity.

We are interested in representations of the solution operator $E(t)$ to the system (1.7) ,

$$u(t,\cdot) = E(t)u_0,$$

in a "basis" $(\boldsymbol{\varphi}_{\mu\nu})$ of $L^2(\mathbb{R}^m)$ as a matrix

$$E(t;\mu,\nu;\mu'\nu') = \langle \boldsymbol{\varphi}_{\mu\nu}, E(t)\boldsymbol{\varphi}_{\mu'\nu'}\rangle. \tag{1.8}$$

In what follows we will take $\boldsymbol{\varphi}_{\mu\nu}(x)$ to be vector-valued curvelets and will explain our choice of notation. Let us remark right away that curvelets do not form orthonormal bases, only tight frames, but that the formula (1.8) is still the relevant object.

1.2.1 A New Form of Multiscale Analysis

Curvelets are waveforms that are highly anisotropic at fine scales, with effective support obeying the parabolic principle $length \approx width^2$. Just as for wavelets, there is both a continuous and a discrete curvelet transform. A curvelet is indexed by three parameters which—adopting a continuous description of the parameter space—are: a scale a, $0 < a < 1$; an orientation θ, $\theta \in [-\pi/2, \pi/2)$ and a location b, $b \in \mathbb{R}^2$. At scale a, the family of curvelets is generated by translation and rotation of a basic element φ_a

$$\varphi_{a,b,\theta}(x) = \varphi_a(R_\theta(x-b)).$$

Here, $\varphi_a(x)$ is some kind of directional wavelet with spatial width $\sim a$ and spatial length $\sim \sqrt{a}$, and with minor axis pointing in the horizontal direction

$$\varphi_a(x) \approx \varphi(D_a x), \quad D_a = \begin{pmatrix} 1/a & 0 \\ 0 & 1/\sqrt{a} \end{pmatrix};$$

D_a is a parabolic scaling matrix, R_θ is a rotation by θ radians. The approximate equality sign indicates that $\varphi_a(D_a^{-1}x)$ may depend on a, but in a non-essential manner.

An important property is that curvelets obey the principle of harmonic of analysis which says that it is possible to analyze and reconstruct an arbitrary function $f(x_1, x_2)$ as a superposition of such templates. It is possible to construct tight frames of curvelets and one can, indeed, easily expand an arbitrary function $f(x_1, x_2)$ as a series of curvelets, much like in an orthonormal basis. Continuing at an informal level of exposition, there is a sampling of the space (a, b, θ)

$$a_j = 2^{-j}, \quad \theta_{j,\ell} = 2\pi\ell \cdot 2^{-\lfloor j/2 \rfloor}, \quad R_{\theta_{j,\ell}} b_k^{(j,\ell)} = (k_1 2^{-j}, k_2 2^{-j/2}),$$

such that with μ indexing the triples $(a_j, \theta_{j,\ell}, b_k^{(j,\ell)})$ the collection φ_μ is a tight frame:

$$f = \sum_\mu \langle f, \varphi_\mu \rangle \varphi_\mu, \quad \|f\|_2^2 = \sum_\mu |\langle f, \varphi_\mu \rangle|^2. \tag{1.9}$$

(Note that these formulae allow us to analyze and synthesize arbitrary functions in $L^2(\mathbb{R}^2)$ as a superposition of curvelets in a stable and concrete way.)

As we have seen, a curvelet is well localized in space but it is also well localized in frequency. Recall that a given scale, curvelets φ_μ are obtained by applying shifts and rotations to a "mother" curvelet $\varphi_{j,0,0}$. In the frequency domain then

$$\hat{\varphi}_{j,0,0}(\xi) = 2^{-3j/4} W(2^{-j}|\xi|)\, V(2^{\lfloor j/2 \rfloor}\theta).$$

Here, W, V are smooth windows compactly supported near the intervals $[1, 2]$ and $[-1/2, 1/2]$ respectively. Whereas in the spatial domain curvelets live near an oriented rectangle R of length $2^{-j/2}$ and width 2^{-j}, in the frequency domain, they are located in a parabolic wedge of length 2^j and width $2^{j/2}$, and whose orientation is orthogonal to that of R. The joint localization in both space and frequency allows us to think about curvelets as occupying a "Heisenberg cell" in phase-space with parabolic scaling in both domains. Figure 1.1 offers a schematic representation of this joint localization. As we shall see, this microlocal behavior is key to understanding the properties of curvelet propagation. Additional details are given in Section 2.2.

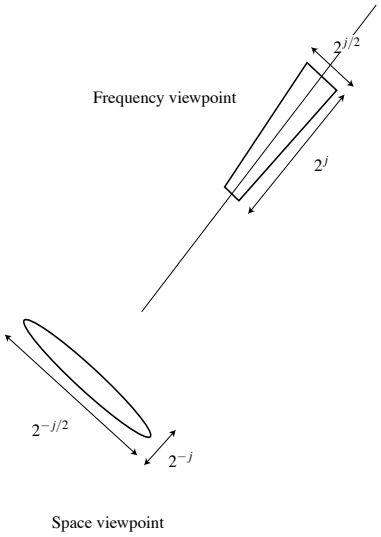

Figure 1.1: Schematic representation of the support of a curvelet in both space and frequency. In the spatial domain, a curvelet has an envelope strongly aligned along a specified 'ridge' while in the frequency domain, it is supported near a box whose orientation is aligned with the codirection of the ridge.

1.2.2 Curvelets and Geometrical Optics

A hyperbolic system can typically be considered in the approximation of high frequency waves, also known as *geometrical optics*. In order to best describe our main result, it is perhaps suitable first to exhibit the connections between curvelets and geometrical optics. In that setting it is not necessary to describe the dynamics in terms of the wavefield $u(t,x)$. Only its prominent features are studied: wave fronts, or equivalently rays. The latter are trajectories $(x(t),\xi(t))$ in phase-space $\mathbb{R}^2 \times \mathbb{R}^2$, and are the solutions to the m Hamiltonian flows (indexed by ν)

$$\begin{cases} \dot{x}(t) &= \nabla_\xi \lambda^0_\nu(x,\xi), & x(0) &= x_0, \\ \dot{\xi}(t) &= -\nabla_x \lambda^0_\nu(x,\xi), & \xi(0) &= \xi_0. \end{cases} \tag{1.10}$$

The system (1.10) is also called the *bicharacteristic flow* and the rays $(x(t),\xi(t))$ the *bicharacteristics*. It is a little bit more general than (1.5) because we are considering general hyperbolic equations. In the above expression, the $\lambda^0_\nu(x,\xi)$ are the eigenvalues of the dispersion matrix

$$a^0(x,\xi) = \sum_k A_k(x)\xi_j. \tag{1.11}$$

(Note that $a^0(x,\xi)$ is the principal symbol of (1.7).) As we mentioned earlier, it is well known that the Hamiltonian equations describe the evolution of the wavefront set of the solution [51].

We are now in a position to qualitatively describe the behavior of the wave propagation operator $E(t)$ acting on a curvelet φ_μ. However, we first need to introduce a notion of vector-valued curvelet since $E(t)$ is acting on vector fields. Let $r^0_\nu(x,\xi)$ be the eigenvector of the dispersion matrix associated with the eigenvalue $\lambda^0_\nu(x,\xi)$. We then define hypercurvelets by

$$\boldsymbol{\varphi}^{(0)}_{\mu\nu}(x) = \frac{1}{(2\pi)^2}\int e^{ix\cdot\xi} r^0_\nu(x,\xi)\hat{\varphi}_\mu(\xi)\,d\xi. \tag{1.12}$$

Later in this section, we will motivate this special choice but for now simply observe that $\boldsymbol{\varphi}^{(0)}_{\mu\nu}$ is a vector-valued waveform.

Consider then the solution to the wave equation $\boldsymbol{\varphi}^{(0)}_{\mu\nu}(t,x)$ with initial value $\boldsymbol{\varphi}^{(0)}_{\mu\nu}(x)$. Our claim is as follows:

the wave group maps each hypercurvelet onto another curvelet-like waveform whose location and orientation are obtained from the corresponding Hamiltonian flow.

To examine this claim, let (x_μ,ξ_μ) be the center of $\boldsymbol{\varphi}^{(0)}_{\mu\nu}$ in phase-space and define the rotation matrix $U(t)$ by

$$U(t)\frac{\xi(t)}{|\xi(t)|} = \frac{\xi_\mu}{|\xi_\mu|}$$

where $(x(t),\xi(t))$ is the solution to (1.10) with initial condition (x_μ,ξ_μ). Our claim says that the solution to the wave equation nearly follows the dynamics of the reduced Hamiltonian flow, i.e.,

$$\boldsymbol{\varphi}^{(0)}_{\mu\nu}(t,x) = \tilde{\boldsymbol{\varphi}}^{(0)}_{\mu\nu}(U_\mu(t)(x-x_\mu(t))+x_\mu). \tag{1.13}$$

We will show in Chapter 2 that the waveform $\tilde{\boldsymbol{\varphi}}^{(0)}_{\mu\nu}$ has the same strong spatial and frequency localization properties as the initial curvelet $\boldsymbol{\varphi}^{(0)}_{\mu\nu}$ itself. For an illustration, see Figure 1.2.

We now return to the interpretation of a hypercurvelet. Suppose that r^0_ν only depends on ξ as in the case of the acoustic system (1.1)

$$r^0_0(\xi) = \begin{pmatrix} \xi^\perp/|\xi| \\ 0 \end{pmatrix}, \qquad r^0_\pm(\xi) = \frac{1}{\sqrt{2}}\begin{pmatrix} \pm\xi/|\xi| \\ 1 \end{pmatrix}.$$

(Here and below, $\xi^\perp$ denotes the vector obtained from ξ after applying a rotation by 90

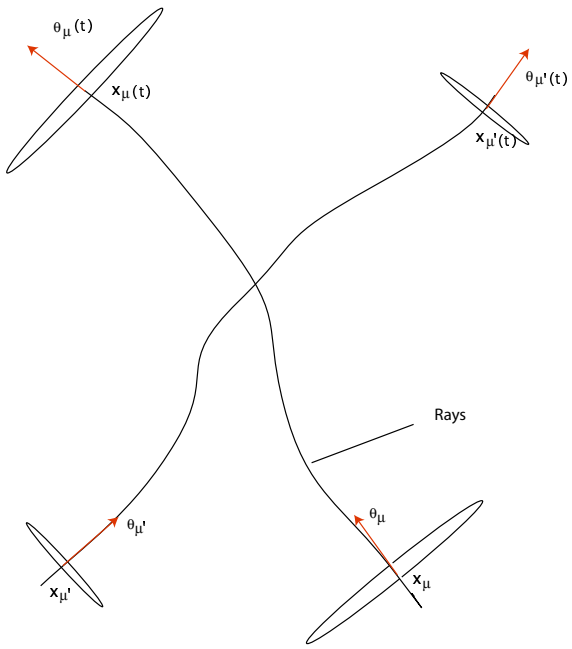

Figure 1.2: Schematic representation of the action of the wave group on a hypercurvelet. The new positions and orientations are given by the Hamiltonian flow. The two waveforms at time 0 and t are not quite the same although they have very similar profiles.

degrees). In this special case, we see that the hypercurvelet is obtained by multiplying—in the frequency domain—a scalar-valued curvelet with the eigenvectors of the dispersion matrix

$$\hat{\boldsymbol{\varphi}}^{(0)}_{\mu\nu}(\xi) = r^0_\nu(\xi)\hat{\varphi}_\mu(\xi), \quad \nu \in \{+,-,0\}.$$

This is useful for the curvelet $\hat{\boldsymbol{\varphi}}^{(0)}_{\mu\nu}$ will essentially follow only one flow, namely, the νth flow. Suppose we had started, instead, with an initial value of the form $\boldsymbol{\varphi}_{\mu\nu} = \varphi_\mu \mathbf{e}_\nu$, where $\mathbf{e}_\nu$ is the canonical basis of $\mathbb{R}^3$, say. Then our curvelet would have interacted with the three eigenvectors of the dispersion matrix, and would have "split" and followed the three distinct flows. By forcing $\hat{\boldsymbol{\varphi}}^{(0)}_{\mu\nu}(\xi)$ to be aligned with $r^0_\nu(\xi)$, we essentially removed the components associated with the other flows. In the general case (1.12), we build hypercurvelets by applying R^0_ν, which is now a pseudodifferential operator with symbol $r^0_\nu(x,\xi)$, mapping scalars to m-dimensional vectors, and independent of time. The effect is, of course, the same.

Note that when r^0_ν is independent of x, hypercurvelets build up a (vector-valued) tight frame; letting $[F,G]$ be the usual inner product over 3D vector fields in $L^2(\mathbb{R}^2)$, the family $(\boldsymbol{\varphi}^{(0)}_{\mu\nu})_{\mu\nu}$ obeys the reconstruction formula

$$u = \sum_{\mu,\nu} [u, \boldsymbol{\varphi}^{(0)}_{\mu\nu}] \boldsymbol{\varphi}^{(0)}_{\mu\nu} \tag{1.14}$$

and the Parseval relation

$$\|u\|^2_{L^2} = \sum_{\mu,\nu} |[u, \boldsymbol{\varphi}^{(0)}_{\mu\nu}]|^2. \tag{1.15}$$

Just as one can decompose a scalar field as a superposition of scalar curvelets, one can analyze and synthesize any wavefield as a superposition of hypercurvelets in a stable and concrete way. For arbitrary $r^0_\nu(x,\xi)$, this is, however, in general not true.

We would like to emphasize that although the Hamilton-Jacobi equations only have solutions for small times, the approximation (1.13) and, more generally, all of our results are valid for all times since the rays (1.10) are always well defined, see Section 1.6 below for a more detailed discussion.

1.2.3 Curvelet and Hyperbolic Systems

The previous section gave a qualitative description of the action of the wave group on a curvelet and we we shall now quantify this fact. The evolution operator $E(t)$ acting on a curvelet $\boldsymbol{\varphi}^{(0)}_{\mu_0\nu_0}$ is of course not exactly another curvelet $\boldsymbol{\varphi}^{(0)}_{\mu_0(t)\nu_0}$ which occurs at a displaced location and orientation. Instead, it is a superposition of curvelets $\sum_{\mu,\mu} \alpha_{\mu\nu}\boldsymbol{\varphi}^{(0)}_{\mu\nu}$ such that

1. the coefficients $(\alpha_{\mu\nu})$ decay nearly exponentially,
2. and the significant coefficients of this expansion are all located at indices (μ, ν) "near" $(\mu_0(t), \nu_0)$. By near, we mean nearby scales, orientations and locations.

To state our key result, we need a notion of distance ω between curvelet indices which will be formally introduced in Section 2.2. Crudely, $\omega(\mu, \mu')$ is small if and only if both curvelets are at roughly the same scale, have similar orientation and are at nearby spatial locations. In the same spirit, the distance $\omega(\mu, \mu')$ increases as the distance between the scale, angular, and location parameters increases.

For each $\mu = (j, k, \ell)$ and $\nu = 1, \ldots, m$, define the vector-valued curvelets

$$\boldsymbol{\varphi}_{\mu\nu} = \mathbf{e}_\nu \varphi_\mu, \tag{1.16}$$

where $\mathbf{e}_\nu$ is the νth canonical basis vector in $\mathbb{R}^m$. The $\boldsymbol{\varphi}_{\mu\nu}$ inherit the tight frame property (1.14)–(1.15). We would like to again remind the reader that these vector-valued curvelets are simpler and different from the hypercurvelets $\boldsymbol{\varphi}^{(0)}_{\mu\nu}$ defined in the previous section. Consider now the representing the operator $E(t)$ in a tight frame of vector-valued curvelets, namely,

$$E(t; \mu, \nu; \mu', \nu') = \langle \boldsymbol{\varphi}_{\mu\nu}, E(t)\boldsymbol{\varphi}_{\mu'\nu'} \rangle. \tag{1.17}$$

We will refer to $E(t; \mu, \nu; \mu', \nu')$ or simply E as the curvelet matrix of $E(t)$, with row index μ, ν and column index μ', ν'. Decompose the initial wavefield $u_0 = \sum_{\mu,\nu} c_{\mu\nu}\boldsymbol{\varphi}_{\mu\nu}$. Then one can express the action of $E(t)$ on u_0 in the curvelet domain as

$$E(t)u_0 = \sum_{\mu\nu} c_{\mu\nu}(t)\boldsymbol{\varphi}_{\mu\nu}, \quad c_{\mu\nu}(t) = \sum_{\mu',\nu'} E(t; \mu', \nu'; \mu, \nu)c_{\mu'\nu'}$$

with convergence in $L^2(\mathbb{R}^2, \mathbb{C}^m)$. In short, the curvelet matrix maps the curvelet coefficients of the initial wavefield $u_0(\cdot)$ into those of the solution $u(t, \cdot)$ at time t.

Theorem 1.1. *Suppose that the coefficients $A_k(x)$ and $B(x)$ of the hyperbolic system are C^∞, and that the multiplicity of the eigenvalues of the dispersion matrix $\sum_k A_k(x)\xi_k$ is constant in x and ξ. Then*

- *The matrix E is sparse. Suppose a is either a row or a column of E, and let $|a|_{(n)}$ be the nth largest entry of the sequence $|a|$, then for each $M > 0$, $|a|_{(n)}$ obeys*

$$|a|_{(n)} \leq C_{tM} \cdot n^{-M}. \tag{1.18}$$

- *The matrix E is well organized. For each $N > 0$, the coefficients obey*

$$|E(t; \mu, \nu; \mu', \nu')| \leq C_{tN} \cdot \sum_{\nu''=1}^{m} \omega(\mu, \mu'_{\nu''}(t))^{-N}. \tag{1.19}$$

  Here $\mu_\nu(t)$ is the curvelet index μ flown along the νth Hamiltonian system.

Both constants C_{tM} and C_{tN} grow in time at most like $C_1 e^{C_2 t}$ for some $C_1, C_2 > 0$ depending on M, resp. N.

In effect, the curvelet matrix of the solution operator resembles a sum of m permutation matrices where m is the order of the hyperbolic system; first, there are significant coefficients along m shifted diagonal, and second, coefficients away from these diagonals decay nearly exponentially; i.e., faster than any negative polynomial. Now just as wavelets provide sparse representations to the solution operators to certain elliptic differential equations, our theorem shows that curvelets provide an optimally sparse representation of solution operators to systems of symmetric hyperbolic equations.

Notice that Theorem 1.1 holds for large times, even when caustics form in the geometrical optics approximation. Caustics are a nonissue in the curvelet domain.

We can also resort to hypercurvelets as defined in the previous section and formulate a related result where the curvelet matrix is sparse around a *single* shifted diagonal. This refinement approximately decouples the evolution into polarized components and will be made precise later.

To grasp the implications of Theorem 1.1, consider the following corollary:

Corollary 1.1. *Consider the truncated operator A_B obtained by keeping $m \cdot B$ elements per row—the B closest to each shifted diagonal in the sense of the pseudo-distance ω. Then the*

truncated matrix obeys

$$\|A - A_B\|_{L^2 \to L^2} \le C_M \cdot B^{-M}, \tag{1.20}$$

for each $M > 0$.

The proof follows from that of Theorem 1.1 by an application of Schur's lemma and is omitted. Hence, whereas the Fourier or wavelet representations are dense, curvelets faithfully model the geometry of wave propagation as only a few terms are needed to represent the action of the wave group accurately.

1.2.4 Generalization: Wave Atoms

It is a natural question to ask whether other waveforms than curvelets would yield comparable sparsity results. The short answer is that the parabolic scaling is essential, allowing only for slight variations on a fixed theme.

Since a complete collection of wave packets $\varphi_\mu(x)$ must "span" all positions and frequencies, we will call it a *phase-space tiling*, with wave packets as tiles. Some tilings are more interesting than others. We say a tiling is universal if it treats democratically all positions and orientations. In that case,

- the geometry of the tiling in space must be Cartesian, or approximately so; and
- the geometry of the tiling in frequency must be polar, or approximately so.

In what follows we limit our discussion to two space variables. This is not an essential restriction.

Universality as above suggests that two parameters should suffice to index a lot of known wave packet architectures: α to index whether the decomposition is "multiscale" ($\alpha = 1$) or not ($\alpha = 0$); and β to indicate whether basis elements should be isotropic ($\beta = \alpha$) or, on the contrary, elongated and anisotropic ($\beta < \alpha$).

In terms of phase-space localization of the wave packets, we will require that

- the essential support of $\varphi_\mu(x)$ be of size $\sim 2^{-\alpha j}$ vs. $2^{-\beta j}$ as scale j, with oscillations of wavelength $\sim 2^{-j}$ tranverse to the ridge; and
- the essential support of $\hat{\varphi}_\mu(\xi)$ be of size $\sim 2^{\alpha j}$ vs. $2^{\beta j}$ as scale j, at a distance $\sim 2^j$ from the origin.

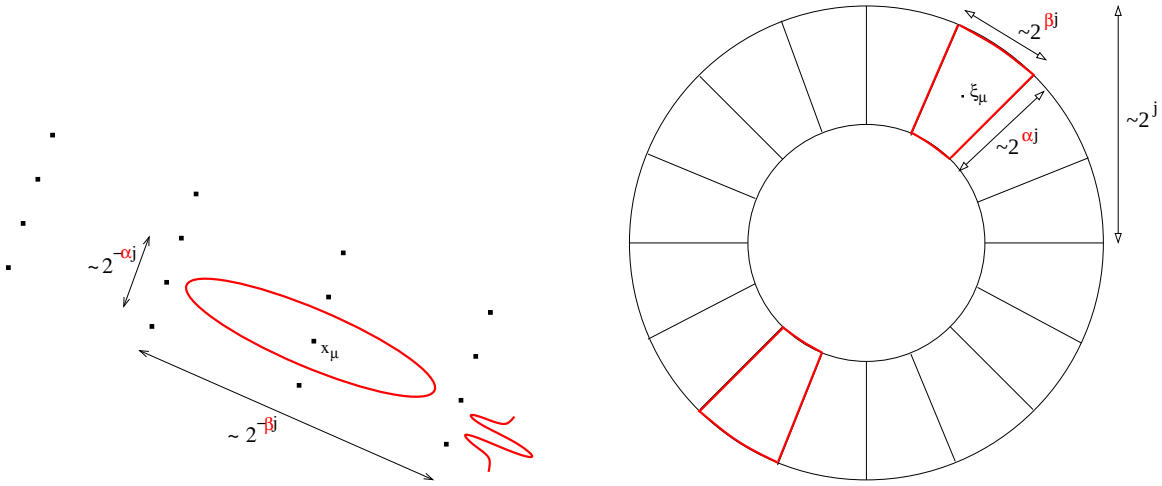

Figure 1.3: Essential support of a wave packet with parameters (α, β), in space (left), and in frequency (right). The parameter α indexes the multiscale nature of the transform, from 0 (uniform) to 1 (dyadic). The parameter β measures the wave packet's directional selectivity, from $\beta = 0$ (best selectivity) to $\beta = 1$ (poor selectivity). Curvelets are the special case $\alpha = 1$, $\beta = 1/2$.

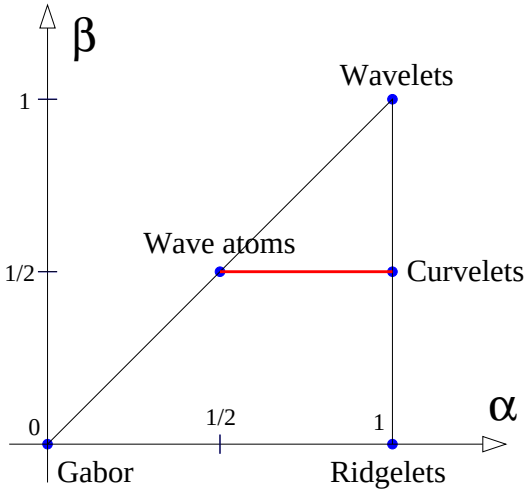

Figure 1.4: Identification of various transforms as (α, β) families of wave packets. The horizontal segment at $\beta = 1/2$ indicates the only wave packet families that yield sparse decompositions of Fourier Integral Operators.

Figure 1.3 summarizes these microlocalization properties.

We hope that a description in terms of α and β will clarify the connections between various transforms of modern harmonic analysis. Curvelets correspond to $\alpha = 1$, $\beta = 1/2$, wavelets are $\alpha = \beta = 1$, ridgelets are $\alpha = 1$, $\beta = 0$, and the Gabor transform is $\alpha = \beta = 0$. The situation is summarized in Figure 1.4.

A careful inspection of the arguments in Chapter 2 shows that the sparsity result, Theorem 1.1, extends to the whole segment $\beta = 1/2$ and $1/2 \le \alpha \le 1$ in the (α, β) plane.

In short, other scalings than the parabolic scaling do not work. In Chapter 2, section 2.3 we argue why that is the case: the parabolic scaling is the right trade-off between directionality and spatial localization for which *dispersion* of waves and *distortion* of wave fronts are comparable and both small.

We term the other endpoint, corresponding to $\alpha = \beta = 1/2$, *wave atoms*, to stress their unique relation to the wave equation. A precise definition of wave atoms will be given in Chapter 4, but let us observe for now that they have an isotropic aspect ratio $\sim 2^{-j/2} \times 2^{-j/2}$ in space, with oscillations of wavelength $\sim 2^{-j}$ in the codirection ξ_μ. Each wave atom is like a train of "stacked curvelets" along ξ_μ.[5]

Note that the range of possible transforms in Figure 1.4 could presumably extend beyond the triangle shown—the horizontal segment indicating sparse FIO, on the other hand, does not. All the transforms within the triangle can be realized as tight frames of $L^2(\mathbb{R}^2)$, in the spirit of the curvelet or wave atom construction that we present next.

As we will see later, wave atoms are sometimes more adequate than curvelets for numerical simulations of wave equations because of their low separation rank.

We now turn our attention to the problem of implementing curvelet and wave atom expansions as digital transforms.

1.3 A Fast Discrete Curvelet Transform

As mentioned earlier, curvelets offer fundamentally improved sparsity rates for functions with discontinuities along curves [20]. As such, their application to wave equations is only one facet of their versatility. A promising potential in application areas such as image

[5] Wave atoms are in a sense simpler than curvelets so we prefer avoiding terming them curvelet packets. Moreover, it seems that wavelet packets refer to the line $\alpha = \beta$, and ridgelet packets to the whole triangle in Figure 1.4 [40] (as well as all other hybrid architectures). We prefer not to add to the confusion.

processing, data analysis, scientific computing, and, in particular, seismic imaging, clearly lies ahead.

To realize this potential though, and deploy this technology to a wide range of problems, one would need a fast and accurate discrete curvelet transform operating on digital data.

1.3.1 Curvelets

Soon after the introduction of curvelets, researchers developed numerical algorithms for their implementation [76, 33], and scientists have started to report on a series of practical successes, see [78, 77] for example. Now these implementations are based on the original construction [18] which uses a preprocessing step involving a special partitioning of phase-space followed by the ridgelet transform [13, 17] which is applied to blocks of data that are well localized in space and frequency.

In 2003 Emmanuel Candès and David Donoho proposed a simplified implementation of second generation curvelets directly in the frequency plane, that relied on interpolation by means of the Unequispaced FFT (USFFT)[16]. Applications can be found at least in [22, 48, 35].

We propose an improved algorithm for second generation curvelets based on the *wrapping* of Fourier samples, instead of interpolation. Wrapping is a simple strategy that results in an equally faithful but faster transform than the USFFT version. It also allows to make the digital transform an *isometry*, up to round-off errors of the order of 10^{-15} in double precision. As a result, the inversion algorithm is greatly simplified, more accurate and simply consists of applying the adjoint transform. Efforts in designing the USFFT and wrapping versions have been grouped and resulted in a single publication, [16].

Since then, it seems that curvelets via wrapping have become the implementation of choice in many applications [49, 47].

We expand on curvelets via wrapping in Chapter 3. Our claims are as follows:

- The complexity of both the forward and inverse transforms is $O(N^2 \log N)$ on N-by-N grids. In practice, applying the transform takes about 5 to 10 times the work of the FFT on the same grid.

- The accuracy of reconstruction is comparable to the machine epsilon.

- The Riesz representers of the transform are as faithful to continuous curvelets as the

grid allows. In particular, we exploit the Shannon sampling theory which says that computing inner products involving bandlimited functions can be done *exactly* on a grid.

Numerical examples are given as well. We show in particular that curvelets via wrapping allow a spectacular gain of 7dB over translation-invariant wavelet thresholding, for denoising of synthetic seismic data.

An extension to three dimensions has been worked out, mostly by Lexing Ying [91]. The curvelet code has been turned into the toolbox CurveLab, and can be downloaded from http://www.curvelet.org.

1.3.2 Wave Atoms

Wave atoms can be implemented using the wrapping strategy in the frequency plane, along the same line of thought as curvelets.

The search for a low redundancy transform is however complicated by the wavelet packet curse, a well documented phenomenon that filterbank ideas provide provably suboptimal time-frequency localization. Our implementation bypasses this obstruction by designing basis functions directly in the frequency plane, with the necessary cancellations properties for numerical tightness (isometry). The ideas involved have roots in harmonic analysis constructions like [88].

We obtain a fast $O(N^2 \log N)$ transform, isometric up to round-off errors, and invertible with inversion algorithm of the same complexity. Wave atoms have redundancy 2, i.e., there are twice more wave atom coefficients than samples on the Cartesian grid. See Chapter 4, Section 4.1.3 for more details.

1.4 Wave Atoms and Time Upscaling of Wave Equations

Typical numerical methods for the wave equation, say in the periodic square $[0,1]^2$ with initial conditions on a N-by-N grid, consist in evolving the solution using small time-steps Δt constrained by the CFL condition,

$$\Delta t < \frac{1}{c_{\max} N},$$

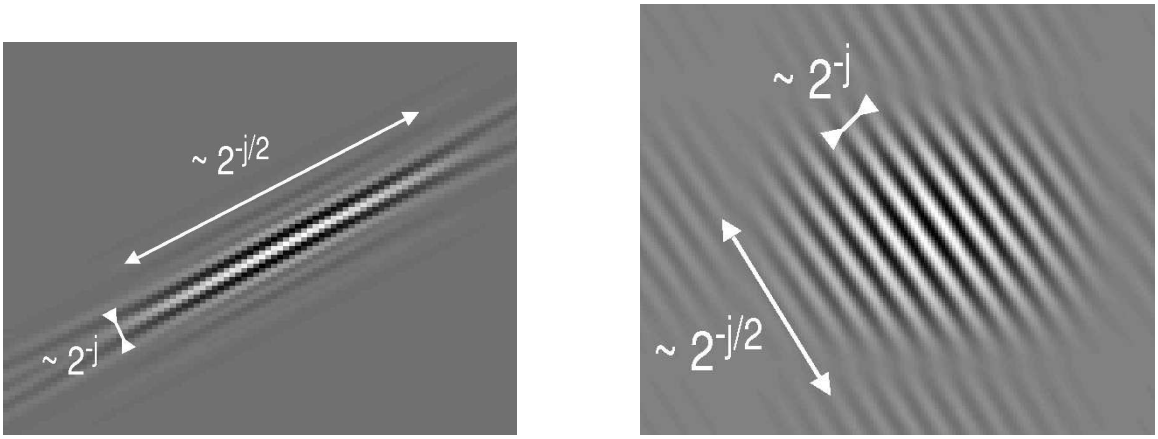

Figure 1.5: A curvelet (left) and a wave atom (right). They are examples of Riesz representers of the digital transforms we propose in this thesis. The annotations remind the reader of the essential parabolic scaling properties: curvelets obey width = length^2 and wave atoms obey wavelength = diameter^2.

where $c_{\max} = \max_x c(x)$. When the wave equation is written as a system $u_t = Au$, the Euler explicit time discretization heuristically converges because

$$u(t) = e^{TA}u_0 \simeq (I + \Delta t\, A)^{\frac{T}{\Delta t}} u_0. \tag{1.21}$$

The results of Sections 1.2 and 1.3 suggest that it may become preferable to solve the same equation by constructing *the full Green's function in compressed form*, and dispensing with the CFL timestepping.

1.4.1 Time Upscaling

We call *time upscaling* the possibility of building a representation of the Green's function $e^{\tau A}$ up to some time τ larger than the CFL timestep Δt, yet smaller than the time T up to which the wave equation needs to be solved. The solution at $t = T$ can be obtained by performing $\frac{T}{\tau}$ "upscaled" time steps, consisting of repeated applications of the Green's function:

$$e^{TA}u_0 = \left(e^{\tau A}\right)^{\frac{T}{\tau}} u_0.$$

The question of representation is left implicit in the above equation. In the sequel we will use a tight frame of curvelets or wave atoms; by writing down A we actually mean the infinite matrix $\langle \boldsymbol{\varphi}_{\mu\nu}, A\boldsymbol{\varphi}_{\mu'\nu'} \rangle$ as in Section 1.2. In that case, u_0 stands for the vector of

coefficients $\langle u_0, \varphi_{\mu',\nu'} \rangle$.

It is important to understand the sense in which a numerical method could qualify as "beating the CFL condition". After all, one could discretize u_0 by finite differences, group small time steps two by two in (1.21), use $(I + \Delta t\, A)^2$ as propagator and declare that the new time step is $2\Delta t$. This operation of course does not qualify as time upscaling, because the matrix representation of the propagator fills up to compensate the larger time step, so that no overall simplification occurs. Progress is achieved only if a representation can be found in which the Green's function *stays* simple, even for times greater than Δt.

On the other hand, perfect time upscaling would be obtained in the basis of eigenfunctions of A. In that case $e^{tA} = \sum_j e^{t\lambda_j} P_j$ is a diagonal operation in each eigenspace with projector P_j. There is, at present, no known fast numerical procedure to compute the eigendecomposition of A in compressed form, let alone expand u_0 in eigenfunctions. Although curvelets and wave atoms are not eigenfunctions, they each offer a *fixed* frame of $L^2([0,1]^2)$ with reliable expansion algorithms and good sparsity properties.

1.4.2 Repeated Squaring

In the spirit of [66] and [37], we form the matrix exponential $e^{\tau A}$ by *repeated squaring* from a small time approximation. Let $t_n = 2^n \Delta t$ for some small Δt, and assume $\tau = t_{n^*}$ for some n^*. Then the basic relation underlying our algorithm is the time-doubling group property

$$e^{t_{n+1}A} = \left(e^{t_n A}\right)^2.$$

As mentioned earlier, this equation should be understood in a tight frame of curvelet or wave atoms.

Sparsity needs to be imposed by an adequate *truncation* step after each time doubling. As we saw in Section 1.2, the large matrix elements occur near two shifted diagonals defined from the flows $\mu'_{\pm}(t)$. Let us call B the desired band size, such that the significant matrix elements live within

$$\omega(\mu, \mu'_{\pm}(t)) \le C_B, \tag{1.22}$$

where C_B is a constant depending on B. Elements outside of those shifted band diagonals should not be accessed or computed at all.

Prediction of the location of the shifted diagonals for $O(1)$ times is not a priori obvious.

We believe the Phase Flow Method (PFM) is perfectly suited for this task [90]. This new method is an important improvement over raytracing which allows to compute a multitude of rays at once. PFM is an interpolation-based repeated squaring strategy to compute the whole *phase flow*, that is, the diffeomorphism of phase-space generated by the Hamiltonian ODE system.

We believe that this simple repeated squaring procedure has near-optimal asymptotic complexity, in the sense that it requires

$$C_{\epsilon,\delta} N^{2+\delta} \tag{1.23}$$

operations to build the propagator on an N-by-N grid, for a resulting ℓ^2 accuracy ϵ on fairly oscillatory initial data, and for arbitrarily small $\delta > 0$. Notice that reading the initial data already takes N^2 operations. The full repeated squaring algorithm can be found in Section 4.2.1.

Heuristically, the complexity result (1.23) follows directly from the compression result (1.20): multiplication of sparse matrices with size N^2 and band size B has complexity $O(B^2 N^2)$. Justifying (1.23) would mean showing that $B = O(N^{\delta/2})$ for small δ suffices to control the error from successive thresholdings and repeated squarings of matrices.

1.4.3 The Separated Wave Atom Representation

To our knowledge, the complexity claim (1.23) for the repeated squaring would be the first to break the asymptotic $O(N^3 \log N)$ bottleneck of standard methods in two dimensions, and by a wide margin. In the spirit of spectral methods, universally good accuracy over oscillatory initial conditions is a result of discretizing differential operators in the Fourier domain. These encouraging result shows that wave packet analysis brings fundamentally new insights into the numerical analysis of wave equations.

Yet, the repeated squaring algorithm as introduced above does not perform as expected, regardless of whether curvelets or wave atoms are used. A typical band size B to obtain ℓ^2 accuracy $\simeq 10^{-2}$ in (1.20) would be $B \simeq 500$. As a result, storing the compressed Green's function on grids larger than 128-by-128 requires more memory than what most 2006 desktop computers can offer (2 to 4 Gb).

Accordingly, we will neither try to formulate the complexity estimate (1.23) as a theorem

in the present thesis, nor address the myriad of (deep and interesting) technicalities involved in its rigorous justification. Instead, we prefer trying to understand how to improve on the algorithm itself.

Asymptotic estimates like (1.20) and (1.23) are probably valid, but with large constants. In two space dimensions these large constants makes sense if we observe that B is the total number of elements inside a ball in four-dimensional phase-space, as in equation (1.22), hence the relation $B \simeq C_B^4$. If $C_B \simeq 5$ elements define a decent neighborhood in phase-space, then $B \simeq 625$.

The message of this section is that the curse of phase-space dimensionality can be overcome with an adequate *separation of variables* strategy in the wave atom frame.

In the notations of Section 1.2.4, consider a tight frame of wave atoms $\varphi_\mu(x)$, with $\mu = (j, \mathbf{m}, \mathbf{n})$. We recall from Section 1.2.3 that the wave atom representation of $E(t)$ is the (infinite) matrix

$$E(t; \mu, \nu; \mu', \nu') = \langle E(t)\varphi_{\mu'}\mathbf{e}_{\nu'}, \varphi_\mu \mathbf{e}_\nu \rangle.$$

where $\mathbf{e}_\nu$ are the canonical basis vectors in $\mathbb{R}^m$. In the above matrix, consider the submatrix left after fixing ν, ν' and the wave vectors $(j, \mathbf{m})$ and $(j', \mathbf{m}')$. The remaining indices are those of the position vectors $\mathbf{n} = (n_1, n_2)$ and $\mathbf{n}' = (n_1', n_2')$. The *separated wave atom representation* is obtained by seeking a low-rank approximation corresponding to separation of the spatial indices along x_1 vs. x_2,

$$E(t; j, \mathbf{m}, \mathbf{n}, \nu; j', \mathbf{m}', \mathbf{n}', \nu') = \sum_{k=1}^{r} \sigma_k u^k_{n_1, n_1'} v^k_{n_2, n_2'} + O(\epsilon),$$

where u^k and v^k have been normalized to unit ℓ_2 norm. Of course u^k and v^k depend on $j, \mathbf{m}, \nu; j', \mathbf{m}', \nu'$. The most efficient such decomposition, in the sense that the ℓ^2 norm of the residual is minimized for fixed r, is the singular value decomposition (SVD) of the block $(j, \mathbf{m}, \nu; j', \mathbf{m}', \nu')$ after *reorganization* of the matrix elements to make the row and column indices $(n_1, n_1'; n_2, n_2')$ instead of $(n_1, n_2; n_1', n_2')$.

Conversion from the standard to the separated wave atom representation, as an SVD factorization of the reorganized submatrix, is however never done in practice. Instead, we modify the repeated squaring strategy so that all computations are done on separated components without ever forming the standard submatrix. We explain in Chapter 4 how both initialization and matrix multiplication can be realized in this context, using small

QR and SVD decompositions.

The composition rules we developed for wave atom submatrices are remindful of the calculus of H-matrices [46]. A similar partitioned low rank representation was also used in [9]. In fact, high-dimensional numerical analysis using separated representations is a promising emerging idea, see [8] and citations thereof.

1.4.4 Complexity

The separated wave atom scheme performs much better than the standard repeated squaring, both in terms of memory and time savings—hence feasibility on larger grids. It even competes with the standard pseudospectral method[6] in regimes where a given wave equation should be solved several times with different initial conditions. Forming the Green's function should be seen as a precomputation that can be amortized over the several runs. For instance in Chapter 4, Section 4.3, we take for $c(x)$ a smooth wave guide and observe that about 500 runs is enough to amortize the precomputation. Upscaled timestepping alone runs 5 to 10 times faster than a pseudospectral method, see Section 4.3.

Complexity of the separated wave atom scheme is very well understood. We give precise estimates of ϵ-separation ranks (r in equation (4.12)) as a function of the upscaled time step τ, the scale j and accuracy level ϵ. The resulting number of operations for repeated squaring (RS) and upscaled timestepping (UTS) are reported in Section 4.2.3. Although not optimal anymore, estimates for UTS still beat the $O(N^3 \log N)$ bottleneck in a variety of physically interesting situations.

The methods of proof of rank estimates rely on understanding the information content of oscillatory functions in high dimensions—or their Fourier dual, functions with singularities—and could be of independent interest in numerical analysis.

Last but not least, *sparsity of the solution wavefield in wave atoms directly translates into complexity gains* for the upscaled timestepping. If the initial condition can be accurately represented using a fraction $\rho < 1$ of all wave atoms, then applying the Green's function in wave atoms only requires considering a fraction ρ of all rows. For instance, we can show that "bandlimited wavefronts" remain so in time and satisfy $\rho = O(\frac{1}{\sqrt{N}})$.

Hence we see that wave atoms, or curvelets, provide the unique opportunity for having a

[6]The pseudospectral method is a split-step timestepping where multiplications by $c(x)$ are done in x, and differentiations are done in the Fourier domain. Periodic boundary conditions implicitly follow from using the FFT.

representation giving enhanced sparsity of wave groups, and *simultaneously* of the solution space. As we alluded to earlier, curvelets are ideal for representing wavefront phenomena [21], or objects which display *curve-punctuated smoothness*—smoothness except for discontinuity along a general curve with bounded curvature [18, 20]. We believe that this joint sparsity property will eventually be of great practical significance for applications in fields which are great consumers of these mathematical models, e.g., seismic imaging.

1.5 Credits

Chapter 2 is joint work with Emmanuel Candès, for the most part published in [14, 15]. In particular, most of the introductory sections were written by Emmanuel, as well as Lemma 2.7 and its proof.

Chapter 3 is joint work with Emmanuel Candès and Lexing Ying. It was published, along with some material not reproduced here, in [16]. The bulk of Section 3.4, as well as the C++ translation of the original Matlab code, are due to Lexing. He is also the sole author of the 3D curvelet transform code.

Chapter 4 is joint work with Emmanuel Candès and Lexing Ying. Section 4.3 was written by Lexing. He is also responsible for the majority of programming that went into creating the wave atom solver in Matlab.

The synthetic seismic data in Chapter 3 is courtesy of Eric Verschuur and Felix Herrmann.

Chapter 2

The Curvelet Representation of Wave Propagators is Optimally Sparse

In this chapter we prove Theorem 1.1. We give some background and introduce the proof strategy in Section 2.1 below. Section 2.2 reviews the construction of Curvelets. Section 2.3 gives further heuristic indicating why the sparsity may be expected to hold. Section 2.4 links our main result with properties of FIOs. Section 2.5 proves that FIOs are optimally sparse in scalar curvelet tight frames. Finally, proofs of key estimates supporting our main result are given in Appendix A.

2.1 Background and Strategy

In his seminal paper [57], Lax constructed approximate solution operators to linear and symmetric hyperbolic systems, also known as *parametrices.* He showed that these parametrices are oscillatory integrals in the frequency domain which are commonly referred to as Fourier integral operators (FIO) (the development and study of FIOs is motivated by the connection). An operator T is said to be an FIO if it is of the form

$$Tf(x) = \int e^{i\Phi(x,\xi)}\sigma(x,\xi)\hat{f}(\xi)\,d\xi. \tag{2.1}$$

We suppose the phase function Φ and the amplitude σ obey the following standard assumptions [79]:

- the phase $\Phi(x,\xi)$ is C^∞, homogeneous of degree 1 in ξ, i.e., $\Phi(x,\lambda\xi) = \lambda\Phi(x,\xi)$ for

$\lambda > 0$, and with $\Phi_{x\xi} = \nabla_x \nabla_\xi \Phi$, obeys the nondegeneracy condition

$$|\det \Phi_{x\xi}(x,\xi)| > c > 0, \tag{2.2}$$

uniformly in x and ξ;

- the amplitude σ is a symbol of order m, which means that σ is C^∞, and obeys

$$|\partial_\xi^\alpha \partial_x^\beta \sigma(x,\xi)| \leq C_{\alpha\beta}(1+|\xi|)^{m-|\alpha|}. \tag{2.3}$$

Lax's insight is that the solution of the initial value problem for a variable coefficient hyperbolic system can be well approximated by a superposition of integrals of the form (2.1) with matrix-valued amplitudes of order 0. The phases of these FIOs are those solving the Hamilton-Jacobi equations

$$\partial_t \Phi_\nu + \lambda_\nu^0(x, \nabla_x \Phi_\nu) = 0 \tag{2.4}$$

(compare with (1.3)). Hence, a substantial part of our argument will be about proving that curvelets sparsify FIOs. Now an important aspect of this construction is that this approximation is only valid for small times whereas our theorem is valid for *all* times. The reason is that the solutions to the Eikonal equations (2.4) are not expected to be global in time, because Φ_ν would become multivalued when rays originating from the same point x_0 cross at a later time. This typically happens at cusp points, when caustics start developing. We refer the reader to [43, 89]. Because, we are interested in a statement valid for all times, we need to bootstrap the construction of the FIO parametrix by composing the small time FIO parametrix with itself. Now this creates an additional difficulty. Each parametrix convects a curvelet along m flows, and we see that after each composition, the number of curvelets would be multiplied by m, see Section 4.1 for a proper discussion. This would lead to matrices with poor concentration properties. Therefore, the other part of the argument consists in decoupling the equations so that this phenomenon does not occur. In summary, the general architecture of the proof of Theorem 1.1 is as follows:

- We first decompose the wave-field into m one-way components, i.e., components which essentially travel along only one flow. We show that this decomposition is sparse in

tight frames of curvelets.

- Second, we show that curvelet representations of FIOs are optimally sparse in tight frame of curvelets, a result of independent interest.

2.2 Curvelets

This section briefly introduces tight frames of curvelets, see [20] for more details.

2.2.1 Definition

We work throughout in $\mathbb{R}^2$, with spatial variable x, with ξ a frequency-domain variable, and with r and θ polar coordinates in the frequency-domain. We start with a pair of windows $W(r)$ and $V(t)$, which we will call the "radial window" and "angular window," respectively. These are both smooth, nonnegative and real valued, with W taking positive real arguments and supported on $r \in [1/2, 2]$ and V taking real arguments and supported on $t \in [-1, 1]$. These windows will always obey the admissibility conditions:

$$\sum_{j=-\infty}^{\infty} W^2(2^j r) = 1, \qquad r > 0; \tag{2.5}$$

$$\sum_{\ell=-\infty}^{\infty} V^2(t-\ell) = 1, \qquad t \in \mathbb{R}. \tag{2.6}$$

Now, for each $j \geq j_0$, we introduce the frequency window U_j defined in the Fourier domain by

$$U_j(r,\theta) = 2^{-3j/4} W(2^{-j} r)\, V(\frac{2^{\lfloor j/2 \rfloor}\theta}{2\pi}). \tag{2.7}$$

where $\lfloor j/2 \rfloor$ is the integer part of $j/2$. Thus the support of U_j is a polar "wedge" defined by the support of W and V, the radial and angular windows, applied with scale-dependent window widths in each direction.

To obtain real-valued curvelets, we could work with the symmetrized version of (2.7), namely, $U_j(r,\theta) + U_j(r,\theta+\pi)$.

Define the waveform $\varphi_j(x)$ by means of its Fourier transform $\hat{\varphi}_j(\omega) = U_j(\omega)$ (we abuse notations slightly here by letting $U_j(\omega_1, \omega_2)$ be the window defined in the polar coordinate system by (2.7)). We may think of φ_j as a "mother" curvelet at scale 2^{-j} in the sense that

all curvelets at that scaled are obtained by rotations and translations of φ_j. Introduce

- the equispaced sequence of *rotation angles* $\theta_\ell = 2\pi \cdot 2^{-\lfloor j/2 \rfloor} \cdot \ell$, with $\ell = 0, 1, \ldots$ such that $0 \le \theta_\ell < 2\pi$ (note that the spacing between consecutive angles is scale-dependent),
- and the sequence of *translation parameters* $k = (k_1, k_2) \in \mathbb{Z}^2$.

With these notations, we define curvelets (as function of $x = (x_1, x_2)$) at scale 2^{-j}, orientation $\theta_{j,\ell}$ and position $b_k^{(j,\ell)} = R_{\theta_{j,\ell}}(k_1 \cdot 2^{-j}/\delta_1, k_2 \cdot 2^{-j/2}/\delta_2)$ for some adequate constants δ_1, δ_2 by

$$\varphi_{j,k,\ell}(x) = \varphi_j \left(R_{-\theta_{j,\ell}}(x - b_k^{(j,\ell)}) \right).$$

where R_θ is the rotation by θ radians and R_θ^{-1} its inverse (also its transpose),

$$R_\theta = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}, \qquad R_\theta^{-1} = R_\theta^T = R_{-\theta}.$$

A curvelet coefficient is then simply the inner product between an element $f \in L^2(\mathbb{R}^2)$ and a curvelet $\varphi_{j,\ell,k}$,

$$c(j, \ell, k) := \langle f, \varphi_{j,\ell,k} \rangle = \int_{\mathbb{R}^2} f(x) \overline{\varphi_{j,\ell,k}(x)}\, dx. \tag{2.8}$$

In the sequel it will prove useful to apply Plancherel's theorem and express this inner product as the integral over the frequency plane

$$c(j, \ell, k) := \frac{1}{(2\pi)^2} \int \hat{f}(\omega)\, \overline{\hat{\varphi}_{j,\ell,k}(\omega)}\, d\omega = \frac{1}{(2\pi)^2} \int \hat{f}(\omega)\, U_j(R_{\theta_\ell}\omega) e^{i\langle x_k^{(j,\ell)}, \omega \rangle}\, d\omega. \tag{2.9}$$

As in wavelet theory, we also have coarse scale elements. We introduce the low-pass window W_0 obeying

$$|W_0(r)|^2 + \sum_{j \ge 0} |W(2^{-j} r)|^2 = 1,$$

and for $k_1, k_2 \in \mathbb{Z}$, define coarse scale curvelets as

$$\Phi_{j_0,k}(x) = \Phi_{j_0}(x - 2^{-j_0} k), \quad \hat{\Phi}_{j_0}(\xi) = 2^{-j_0} W_0(2^{-j_0}|\xi|).$$

Hence, coarse scale curvelets are nondirectional. The full curvelet transform consists of the fine-scale directional elements $(\varphi_{j,\ell,k})_{j \ge j_0, \ell, k}$ and of the coarse-scale isotropic father wavelets

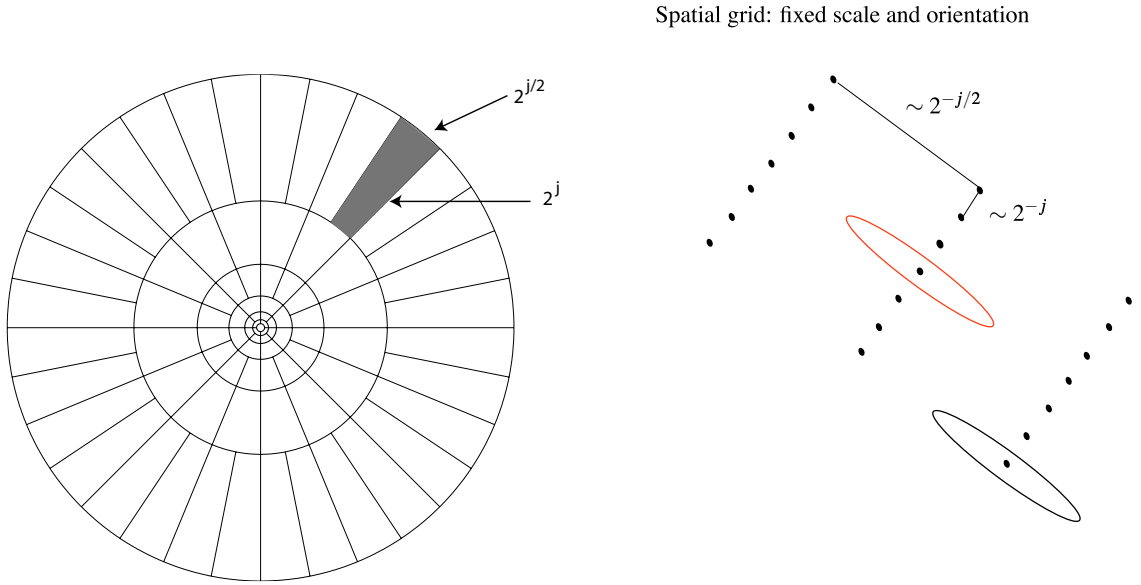

Figure 2.1: Curvelet tiling of phase-space. The figure on the left represents the sampling in the frequency plane, also called second dyadic decomposition (SDD). In the frequency domain, curvelets are supported near a "parabolic" wedge. The shaded area represents such a generic wedge. The figure on the right schematically represents the spatial Cartesian grid associated with a given scale and orientation.

$(\Phi_{j_0,k})_k$. It is the behavior of the fine-scale directional elements that are of interest here.

In the remainder of the chapter, we will use the generic notation $(\varphi_\mu)_{\mu\in M}$ to index the elements of the curvelet tight frame. The dyadic-parabolic subscript μ stands for the triplet (j,k,ℓ). We will also make use of the convenient notations

- $x_\mu = b_k^{(j,\ell)}$ is the center of φ_μ in space.
- $\theta_\mu = \theta_{j,\ell}$ is the orientation of φ_μ with respect to the vertical axis in x.
- $\xi_\mu = (2^j \cos\theta_\mu, 2^j \sin\theta_\mu)$ is the center of $\hat{\varphi}_\mu$ in frequency.
- $e_\mu = \xi_\mu/|\xi_\mu|$ indicates the codirection of φ_μ.

Figure 2.1 summarizes the key components of the construction.

2.2.2 Properties

We now list a few properties of the curvelet transform which will play an important role throughout the remainder of this thesis.

1. **Tight frame**. Much like in an orthonormal basis, we can easily expand an arbitrary function $f(x_1, x_2) \in L^2(\mathbb{R}^2)$ as a series of curvelets: we have a reconstruction formula

$$f = \sum_{\mu} \langle f, \varphi_\mu \rangle \varphi_\mu,$$

with equality holding in an L^2 sense; and a Parseval relation

$$\sum_{\mu} |\langle f, \varphi_\mu \rangle|^2 = \|f\|^2_{L^2(\mathbb{R}^2)}, \quad \forall f \in L^2(\mathbb{R}^2).$$

2. **Parabolic scaling**. The frequency localization of φ_j implies the following spatial structure: $\varphi_j(x)$ is of rapid decay away from a 2^{-j} by $2^{-j/2}$ rectangle with minor axis pointing in the horizontal direction. In short, the effective length and width obey the anisotropy scaling relation

$$length \approx 2^{-j/2}, \quad width \approx 2^{-j} \quad \Rightarrow \quad width \approx length^2. \tag{2.10}$$

3. **Oscillatory behavior**. As is apparent from its definition, $\hat{\varphi}_j$ is actually supported away from the vertical axis $\xi_1 = 0$ but near the horizontal $\xi_2 = 0$ axis. In a nutshell, this says that $\varphi_j(x)$ is oscillatory in the x_1-direction and lowpass in the x_2-direction. Hence, at scale 2^{-j}, a curvelet is a little needle whose envelope is a specified "ridge" of effective length $2^{-j/2}$ and width 2^{-j}, and which displays an oscillatory behavior across the main "ridge."

4. **Vanishing moments**. The curvelet template φ_j is said to have q vanishing moments when

$$\int_{-\infty}^{\infty} \varphi_j(x_1, x_2) x_1^n \, dx_1 = 0, \qquad \text{for all } 0 \le n < q, \text{ for all } x_2. \tag{2.11}$$

The same property of course holds for rotated curvelets when x_1 and x_2 are taken to be the corresponding rotated coordinates. Notice that the integral is taken in the direction perpendicular to the ridge, so counting vanishing moments is a way to quantify the oscillation property mentioned above. In the Fourier domain, (2.11)

becomes a line of zeros with some multiplicity:

$$\frac{\partial^n \hat{\varphi}_j}{\partial \omega_1^n}(0, \omega_2) = 0, \qquad \text{for all } 0 \le n < q, \text{ for all } \omega_2.$$

Curvelets as defined and implemented in this thesis have an infinite number of vanishing moments because they are compactly supported well away from the origin in the frequency plane, as illustrated in Figure 2.1.

5. **Phase-Space Tiling/Sampling**. We can really think about curvelets as Heisenberg tiles of minimum volume in phase-space. In x, the essential support of φ_μ has size $O(2^{-j} \times 2^{-j/2})$. In frequency, the support of $\hat{\varphi}_\mu$ has size $O(2^{j/2} \times 2^j)$. The net volume in phase-space is therefore

$$O(2^{-j} \times 2^{-j/2}) \cdot O(2^{j/2} \times 2^j) = O(1),$$

which is in accordance with the uncertainty principle. The parameters (j, k, ℓ) of the curvelet transform induce a new non-trivial sampling of phase-space, Cartesian in x, polar in ξ, and based on the parabolic scaling.

6. **Complex-valuedness**. Since curvelets do not obey the symmetry $\hat{\varphi}_\mu(-\xi) = \overline{\hat{\varphi}_\mu(\xi)}$, φ_μ is complex-valued. There exists a related construction for real-valued curvelets by simply symmetrizing the construction, see [20]. The complex-valued transform is better adapted to the purpose of this chapter.

2.2.3 Curvelet Molecules

We introduce the notion of *curvelet molecule*; our objective, here, is to encompass under this name a wide collection of systems which share the same essential properties as the curvelets we have just introduced. Our formulation is inspired by the notion of "vaguelettes" in wavelet analysis [65]. Our motivation for introducing this concept is the fact that operators of interest do not map curvelets into curvelets, but rather into these molecules. Note that the terminology "molecule" is somewhat standard in the literature of harmonic analysis [42].

Definition 2.1. *A family of functions $(m_\mu)_\mu$ is said to be a family of curvelet molecules*

with regularity R if (for $j > 0$) they may be expressed as

$$m_\mu(x) = 2^{3j/4} a^{(\mu)} \left(D_{2^{-j}} R_{\theta_\mu} x - k' \right),$$

where $k' = (\frac{k_1}{\delta_1}, \frac{k_2}{\delta_2})$ and where for all μ, the $a^{(\mu)}$'s verify the following properties:

- *Smoothness and spatial localization: for each $|\beta| \leq R$, and each $M = 0, 1, 2, \ldots$ there is a constant $C_M > 0$ such that*

$$|\partial_x^\beta a^{(\mu)}(x)| \leq C_M \cdot (1 + |x|)^{-M}. \tag{2.12}$$

- *Nearly vanishing moments: for each $N = 0, 1, \ldots, R$, there is a constant $C_N > 0$ such that*

$$|\hat{a}^{(\mu)}(\xi)| \leq C_N \cdot \min(1, 2^{-j} + |\xi_1| + 2^{-j/2}|\xi_2|)^N. \tag{2.13}$$

Here, the constants may be chosen independently of μ so that the above inequalities hold uniformly over μ. There is of course an obvious modification for the coarse scale molecules which are of the form $a^{(\mu)}(x - k')$ with $a^{(\mu)}$ as in (2.12).

This definition implies a series of useful estimates. For instance, consider $\theta_\mu = 0$ so that R_{θ_μ} is the identity (arbitrary molecules are obtained by rotations). Then, m_μ obeys

$$|m_\mu(x)| \leq C_M \cdot 2^{3j/4} \cdot \left(1 + |2^j x_1 - \frac{k_1}{\delta_1}| + |2^{j/2} x_2 - \frac{k_2}{\delta_2}| \right)^{-M} \tag{2.14}$$

for each $M > 0$ and $|\beta| \leq R$, and similarly for its derivatives

$$|\partial_x^\beta m_\mu(x)| \leq C_M \cdot 2^{3j/4} \cdot 2^{(\beta_1 + \beta_2/2)j} \cdot \left(1 + |2^j x_1 - \frac{k_1}{\delta_1}| + |2^{j/2} x_2 - \frac{k_2}{\delta_2}| \right)^{-M}. \tag{2.15}$$

Another useful property is the almost vanishing moments property which says that in the frequency plane, a molecule is localized near the dyadic corona $\{2^j \leq |\xi| \leq 2^{j+1}\}$; $|\hat{m}_\mu(\xi)|$ obeys

$$|\hat{m}_\mu(\xi)| \leq C_N \cdot 2^{-3j/4} \cdot \min(1, 2^{-j}(1 + |\xi|))^N, \tag{2.16}$$

which is valid for every $N \leq R$, which gives the frequency localization

$$|\hat{m}_\mu(\xi)| \leq C_N \cdot 2^{-3j/4} \cdot |S_\mu(\xi)|^N, \tag{2.17}$$

where for $\mu_0 = (j, 0, 0)$,

$$S_{\mu_0}(\xi) = \min(1, 2^{-j}(1 + |\xi|)) \cdot (1 + |2^{-j}\xi_1| + |2^{-j/2}\xi_2|)^{-1}. \tag{2.18}$$

For arbitrary μ, S_μ is obtained from S_{μ_0} by a simple rotation of angle θ_μ, i.e., $S_{\mu_0}(R_{\theta_\mu}\xi)$. Similar estimates are available for the derivatives of $\hat{\varphi}_\mu$.

In short, a curvelet molecule is a needle whose envelope is supported near a ridge of length about $2^{-j/2}$ and width 2^{-j} and which displays an oscillatory behavior across the ridge. It is easy to show that curvelets as introduced in the previous section are indeed curvelet molecules for arbitrary degrees R of regularity.

2.2.4 Near Orthogonality of Curvelet Molecules

Curvelets are not necessarily orthogonal to each other,[1] but in some sense they are almost orthogonal. As we show below, the inner product between two molecules m_μ and $p_{\mu'}$ decays nearly exponentially as a function of the "distance" between the subscripts μ and μ'.

This notion of distance in phase-space, tailored to curvelet analysis, is to be understood as follows. Given a pair of indices $\mu = (j, k, \ell), \mu' = (j', k', \ell')$, define the *dyadic-parabolic pseudodistance*

$$\omega(\mu, \mu') = 2^{|j-j'|} \cdot \left(1 + \min(2^j, 2^{j'})\, d(\mu, \mu')\right), \tag{2.19}$$

where

$$d(\mu, \mu') = |\theta_\mu - \theta_{\mu'}|^2 + |x_\mu - x_{\mu'}|^2 + |\langle e_\mu, x_\mu - x_{\mu'}\rangle|.$$

Angle differences like $\theta_\mu - \theta_{\mu'}$ are understood modulo π. As introduced earlier, e_μ is the codirection of the first molecule, i.e., $e_\mu = (\cos\theta_\mu, \sin\theta_\mu)$.

The pseudodistance (2.19) is a slight variation on that introduced by Smith [74]. We see that ω increases by at most a constant factor every time the distance between the scale, angular, and location parameters increases. The extension of the definition of ω to arbitrary points (x, ξ) and (x', ξ') is straightforward. Observe that the extra term $|\langle e_\mu, x_\mu - x_{\mu'}\rangle|$ induces a non-Euclidean notion of distance between x_μ and $x_{\mu'}$. The following properties of ω are proved in Appendix A.1. (The notation $A \asymp B$ means that $C_1 \le A/B \le C_2$ for some constants $C_1, C_2 > 0$.)

[1] It is an open problem whether orthobases of curvelets exist or not.

Proposition 2.1. *1. Symmetry:* $\omega(\mu,\mu') \asymp \omega(\mu',\mu)$.

2. Triangle inequality: $d(\mu,\mu') \le C\cdot(d(\mu,\mu'')+d(\mu'',\mu'))$ *for some constant* $C>0$.

3. Composition: for every integer $N>0$, *and some positive constant* C_N

$$\sum_{\mu''} \omega(\mu,\mu'')^{-N}\cdot\omega(\mu'',\mu')^{-N} \le C_N\cdot\omega(\mu,\mu')^{-(N-1)}.$$

4. Invariance under Hamiltonian flows: $\omega(\mu,\mu') \asymp \omega(\mu_\nu(t),\mu'_\nu(t))$.

We can now state the almost orthogonality result

Lemma 2.1. *Let* $(m_\mu)_\mu$ *and* $(p_{\mu'})_{\mu'}$ *be two families of curvelet molecules with regularity* R. *Then for* $j,j'\ge 0$,

$$|\langle m_\mu, p_{\mu'}\rangle| \le C_N\cdot\omega(\mu,\mu')^{-N}. \tag{2.20}$$

for every $N\le f(R)$ *where* $f(R)$ *goes to infinity as* R *goes to infinity.*

Proof. Throughout the proof of (2.20), it will be useful to keep in mind that $A\le C\cdot(1+|B|)^{-M}$ for every $M\le 2M'$ is equivalent to $A\le C\cdot(1+B^2)^{-M}$ for every $M\le M'$. Similarly, if $A\le C\cdot(1+|B_1|)^{-M}$ and $A\le C\cdot(1+|B_2|)^{-M}$ for every $M\le 2M'$, then $A\le C\cdot(1+|B_1|+|B_2|)^{-M}$ for every $M\le M'$. Here and throughout, the constants C may vary from expression to expression.

For notational convenience put $\Delta\theta=\theta_\mu-\theta_{\mu'}$ and $\Delta x = x_\mu - x_{\mu'}$. We abuse notation by letting m_{μ_0} be the molecule $a^{(\mu)}(D_{2^{-j}}R_{\theta_\mu}x)$, i.e., m_{μ_0} is obtained from m_μ by translation so that it is centered near the origin. Put $I_{\mu\mu'}=\langle m_\mu,p_{\mu'}\rangle$. In the frequency domain, $I_{\mu\mu'}$ is given by

$$I_{\mu\mu'} = \frac{1}{(2\pi)^2}\int \hat m_{\mu_0}(\xi)\overline{\hat p_{\mu'_0}(\xi)}\, e^{-i(\Delta x)\cdot\xi}\,d\xi.$$

Put j_0 to be the minimum of j and j'. The Appendix shows that

$$\int |S_{\mu_0}(\xi)\,S_{\mu'_0}(\xi)|^N\,d\xi \le C\cdot 2^{3j/4+3j'/4}\cdot 2^{-|j-j'|N}\cdot\left(1+2^{j_0}|\Delta\theta|^2\right)^{-N}, \tag{2.21}$$

where S_{μ_0} is defined in equation (2.18). Therefore, the frequency localization of the curvelet

molecules (2.17) gives

$$\begin{aligned}\int |\hat{m}_{\mu_0}(\xi)|\,|\hat{p}_{\mu_0'}(\xi)|\,d\xi &\leq C\cdot 2^{-3j/4-3j'/4}\cdot\int |S_{\mu_0}(\xi)\,S_{\mu_0'}(\xi)|^N\,d\xi \\ &\leq C\cdot 2^{-|j-j'|N}\cdot\left(1+2^{j_0}|\Delta\theta|^2\right)^{-N}. \qquad (2.22)\end{aligned}$$

This inequality explains the angular decay. A series of integrations by parts will introduce the spatial decay, as we now show.

The partial derivatives of $\hat{m}_\mu$ obey

$$|\partial_\xi^\alpha \hat{m}_\mu(\xi)| \leq C\cdot 2^{-3j/4}\cdot 2^{-j(\alpha_1+\frac{\alpha_2}{2})}\cdot |S_\mu(\xi)|^N.$$

Put Δ_ξ to be the Laplacian in ξ. Because $\hat{p}_{\mu'}$ is misoriented with respect to e_μ, simple calculations show that

$$\begin{aligned}&|\Delta_\xi \hat{p}_{\mu'}(\xi)| \leq C\cdot 2^{-3j'/4}\cdot 2^{-j'}\cdot |S_{\mu'}(\xi)|^N,\\ &|\frac{\partial^2}{\partial\xi_1^2}\hat{p}_{\mu'}(\xi)| \leq C\cdot 2^{-3j'/4}\cdot (2^{-2j'}+2^{-j'}|\sin(\Delta\theta)|^2)\cdot |S_{\mu'}(\xi)|^N.\end{aligned}$$

Recall that for $t\in[-\pi/2,\pi/2]$, $2/\pi\cdot|t|\leq|\sin t|\leq|t|$, so we may just as well replace $|\sin(\Delta\theta)|$ by $|\Delta\theta|$ in the above inequality. Set

$$L = I - 2^{j_0}\Delta_\xi - \frac{2^{2j_0}}{1+2^{j_0}|\Delta\theta|^2}\frac{\partial^2}{\partial\xi_1^2}.$$

On the one hand, for each k, $L^k(\hat{m}_\mu\overline{\hat{p}_{\mu'}})$ obeys

$$|L^k(\hat{m}_\mu\overline{\hat{p}_{\mu'}})(\xi)| \leq C\cdot 2^{-3j/4-3j'/4}\cdot|S_\mu(\xi)|^N\cdot|S_{\mu'}(\xi)|^N.$$

On the other hand

$$L^k e^{-i(\Delta x)\cdot\xi} = [1+2^{j_0}|\Delta x|^2+\frac{2^{2j_0}}{1+2^{j_0}|\Delta\theta|^2}|\langle e_\mu,\Delta x\rangle|^2]^k e^{-i(\Delta x)\cdot\xi}.$$

Therefore, a few integrations by parts give

$$|I_{\mu\mu'}| \leq C \cdot 2^{-|j-j'|N} \cdot \left(1 + 2^{j_0}|\theta_\mu - \theta_{\mu'}|^2\right)^{-N} \cdot \left(1 + 2^{j_0}|\Delta x|^2 + \frac{2^{2j_0}}{1 + 2^{j_0}|\Delta\theta|^2}|\langle e_\mu, \Delta x\rangle|^2\right)^{-N},$$

and then

$$|I_{\mu\mu'}| \leq C \cdot 2^{-|j-j'|M} \cdot \left(1 + 2^{j_0}(|\Delta\theta|^2 + |\Delta x|^2) + \frac{2^{2j_0}}{1 + 2^{j_0}|\Delta\theta|^2}|\langle e_\mu, \Delta x\rangle|^2\right)^{-N}.$$

One can simplify this expression by noticing that

$$(1 + 2^{j_0}|\Delta\theta|^2) + \frac{2^{2j_0}|\langle e_\mu, \Delta x\rangle|^2}{1 + 2^{j_0}|\Delta\theta|^2} \gtrsim \sqrt{1 + 2^{j_0}|\Delta\theta|^2}\frac{2^{j_0}|\langle e_\mu, \Delta x\rangle|}{\sqrt{1 + 2^{j_0}|\Delta\theta|^2}} = 2^{j_0}|\langle e_\mu, \Delta x\rangle|.$$

This yields equation (2.20) as required. □

Remark. Assume that one of the two terms or both terms are coarse scale molecules, e.g., $p_{\mu'}$, then the decay estimate is of the form

$$|\langle m_\mu, p_{\mu'}\rangle| \leq C \cdot 2^{-jN} \cdot \left(1 + |x_\mu - x_{\mu'}|^2 + |\langle e_\mu, x_\mu - x_{\mu'}\rangle|\right)^{-N}.$$

For instance, if they are both coarse scale molecules, this would give

$$|\langle m_\mu, p_{\mu'}\rangle| \leq C \cdot \left(1 + |x_\mu - x_{\mu'}|\right)^{-N}.$$

The following result is a different expression for the almost-orthogonality, and will be at the heart of the sparsity estimates for FIOs.

Lemma 2.2. *Let $(m_\mu)_\mu$ and $(p_\mu)_\mu$ be two families of curvelet molecules with regularity R. Then for each $p > p^*$,*

$$\sup_\mu \sum_{\mu'} |\langle m_\mu, p_{\mu'}\rangle|^p \leq C_p.$$

Here $p^* \to 0$ as $R \to \infty$. In other words, for $p > p^*$, the matrix $I_{\mu\mu'} = (\langle m_\mu, p_{\mu'}\rangle)_{\mu,\mu'}$ acting on sequences (α_μ) obeys

$$\|I\alpha\|_{\ell_p} \leq C_p \cdot \|\alpha\|_{\ell_p}.$$

Proof. Put as before $j_0 = \min(j, j')$. The appendix shows that

$$\sum_{\mu \in M_{j'}} \left(1 + 2^{j_0}(d(\mu, \mu')\right)^{-Np} \leq C \cdot 2^{2|j-j'|} \tag{2.23}$$

provided that $Np > 2$. We then have

$$\sum_{\mu'} |I_{\mu\mu'}|^p \leq C \cdot \sum_{j' \in \mathbb{Z}} 2^{-2|j-j'|Np} \cdot 2^{2|j-j'|} \leq C_p,$$

provided again that $Np > 2$.

Hence we proved that for $p \leq 1$, I is a bounded operator from ℓ_p to ℓ_p. We can of course interchange the role of the two molecules and obtain

$$\sup_{\mu'} \sum_{\mu} |\langle m_\mu, p_{\mu'} \rangle|^p \leq C_p.$$

For $p = 1$, the above expression says that I is a bounded operator from ℓ_∞ to ℓ_∞. By interpolation, we then conclude that I is a bounded operator from ℓ_p to ℓ_p for every p. □

This completes the discussion of the mathematical properties of curvelets. We can now go back to hyperbolic equations.

2.3 Heuristics

This section explains the organization of the argument underlying the proof of the main result, namely, Theorem 1.1, and gives the main reasons why curvelets are special.

2.3.1 Architecture of the Proof of the Main Result

- *Decoupling into polarized components.* The first step is to decouple the wavefield $u(t, x)$ into m one-way components $f_\nu(t, x)$

$$u(t, x) = \sum_{\nu=1}^{m} R_\nu f_\nu(t, x),$$

where the R_ν are operators mapping scalars to m-dimensional vectors, and independent of time. The f_ν will also be called "polarized" components. This allows

a separate study of the m flows corresponding to the m eigenvalues of the matrix $\sum_{k=1}^{m} A_k(x)\xi_k$. In the event these eigenvalues are simple, the evolution operator $E(t)$ can be decomposed as

$$E(t) = \sum_{\nu=1}^{m} R_\nu e^{-it\Lambda_\nu} L_\nu + \text{ negligible}, \tag{2.24}$$

where the L_ν are operators mapping m-dimensional vectors to scalars and the Λ_ν's are one-way wave operators acting on scalar functions. In effect, each operator $E_\nu(t) = e^{-it\Lambda_\nu}$ convects wave-fronts and other singularities along a separate flow. The "negligible" contribution is a smoothing operator—not necessarily small. The composition operators R_ν and decomposition operators L_ν are provably pseudodifferential operators, see Section 2.4.2.

- *Fourier integral operator parametrix.* We then approximate for small times $t > 0$ each $e^{-it\Lambda_\nu}$, $\nu = 1, \ldots, m$, by an oscillatory integral or Fourier integral operator (FIO) $F_\nu(t)$. Such operators take the form

$$F_\nu(t)f(x) = \int e^{i\Phi_\nu(t,x,\xi)} \sigma_\nu(t,x,\xi) \hat{f}(\xi)\, d\xi,$$

under suitable conditions on the phase function $\Phi_\nu(t,x,\xi)$ and the amplitude $\sigma_\nu(t,x,\xi)$. Again, the identification of the evolution operator $E_\nu(t) = e^{-it\Lambda_\nu}$ with F_ν is valid up to a smoothing and localized additive remainder. The construction of the so-called *parametrix* $F_\nu(t)$ and its properties are detailed in Section 2.4.3.

Historically [57], the construction of an oscillatory integral parametrix did not involve the decoupling into polarized components as a preliminary step. When applied directly to the system (2.25), the construction of the parametrix gives rise to a matrix-valued amplitude $\sigma(t,x,\xi)$ where all the couplings are present. This somewhat simpler setting, however, is not adequate for our purpose. The reason is that we want to bootstrap the construction of a parametrix to *large* times by composing the small time FIO parametrix with itself, $F(nt) = [F(t)]^n$. Without decoupling of the propagation modes, each $E(t)$ or $F(t)$ involves convection of singularities along m families of characteristics or flows. Applying $F(t)$ again, each flow would artificially split into m flows again, yielding m^2 fronts to keep track of. At time $T = nt$, that would be at most

m^n fronts. This flow-splitting situation is not physical and can be avoided by isolating one-way components *before* constructing the parametrix. The correct large-time argument is to consider $E_\nu(nt)$ for small $t > 0$ and large integer n as $[E_\nu(t)]^n$. This expression involves one single flow, indexed by ν.

- *Sparsity of Fourier integral operators.* The core of the proof is found in Section 2.5 and consists in showing that very general FIOs $F(t)$, including the parametrices $F_\nu(t)$, are sparse and well structured when represented in tight frames of (scalar) curvelets φ_μ. The scalar analog of Theorem 1.1 for FIOs is Theorem 2.1—a statement of independent interest. Observe that pseudodifferential operators are a special class of FIOs and, therefore, are equally sparse in a curvelet frame.

Section 2.4.5 assembles key intermediate results and proves Theorem 1.1.

2.3.2 The Parabolic Scaling is Special

Why is the curvelet parabolic scaling the only correct way to scale a family of wave packets to sparsely represent wave groups? In analogy with the discussion in Section 1.2.4, assume for a moment that the curvelet scaling $width \approx length^2$ is replaced with the more general power-law

$$width \approx length^{\frac{1}{\beta}}, \quad 1 \leq \frac{1}{\beta} \leq \infty,$$

and that one has available a tight frame φ_μ of "β-wave-packets" (we have put $\alpha = 1$ in the notations of Section 1.2.4.) For example, $\beta = 1$ would correspond to wavelets and $\beta = 0$ to ridgelets [13].

Consider a wave packet $\varphi_\mu(x)$ centered around x_μ in space and ξ_μ in frequency. The action of a Fourier integral operator on this wave packet can be viewed as the composition of two transformations, (1) non-rigid convection along the Hamiltonian flow due to the phase factor $\Phi(x, t, \xi)$ (or more precisely its linearization $\xi \cdot \nabla_\xi \Phi(t, x, \xi_\mu)$ around ξ_μ) and (2) microlocal dispersion due to the remainder after linearization and the amplitude $\sigma(t, x, \xi)$. Depending upon the size of the essential support in phase-space (controlled by the value of β), these two transformations may leave the shape of the waveform nearly invariant, or not. We now argue that the curvelet parabolic scaling, $\beta = 1/2$, offers the correct compromise.

1. **Spatial localization**. For simplicity, suppose that one can model the convective

effect by a smooth diffeomorphism $g(x)$, so that a wave-packet $\varphi_\mu(x)$ is effectively mapped into $\varphi_\mu(g(x))$. If we Taylor expand $g(x)$ around $y_\mu = g^{-1}(x_\mu)$, where x_μ is the center of $\phi_\mu(x)$, we obtain

$$g(x) = x_\mu + (x - x_\mu)g'(y_\mu) + O((x - x_\mu)^2).$$

The first two terms induce an essentially rigid motion, while the remainder is responsible for deforming the waveform. The requirement for optimal sparsity, as it turns out, is that the extent of the deformation should not exceed the width of the wave packet. In the case of curvelets, this imposes the correct condition for $\varphi_\mu(g(x))$ to remain a "molecule" in the sense defined earlier; see how equation (2.71) combines with the molecule estimate (2.14).

If the spatial width is of the order of $a = 2^{-j}$, then the wave packet should essentially be supported in a region obeying $(x - x_\mu) \sim 2^{-j/2}$. This is satisfied if and only if $1/2 \leq \beta \leq 1$. In short, *any scaling more isotropic than the parabolic scaling works.*

2. **Frequency localization**. Dispersive effects are already present in the wave equation with constant velocity $c = 1$,

$$\frac{\partial^2 u}{\partial t^2} = \Delta u,$$

with initial conditions $u(x, 0) = u_0(x)$, $\frac{\partial u}{\partial t}(x, 0) = u_1(x)$. In the Fourier domain, the solution is given by

$$\hat{u}(t, \xi) = \cos(|\xi|t)\, \hat{u}_0(\xi) + \frac{\sin(|\xi|t)}{|\xi|}\, \hat{u}_1(\xi).$$

These multipliers are of course associated with the phases $\Phi_\pm(t, x, \xi) = x \cdot \xi \pm t|\xi|$ (express sine and cosine in terms of complex exponentials). Linearize $\Phi_\pm$ around ξ_μ, the center of $\hat{\varphi}_\mu$, and obtain

$$\Phi_\pm(t, x, \xi) = (x \pm t\mathbf{e}_\mu) \cdot \xi \pm t(|\xi| - \xi \cdot \mathbf{e}_\mu),$$

where $\mathbf{e}_\mu = \frac{\xi_\mu}{|\xi_\mu|}$. The first term is responsible for convection as before while the second is responsible for dispersion (transverse to the oscillations of the wave packet). Again, we must invoke more sophisticated arguments to see that to achieve sparsity,

one needs $\delta(\xi) = |\xi| - \xi \cdot \mathbf{e}_\mu$ to be uniformly bounded over the frequency support of the wave packet $\hat{\varphi}_\mu$ as to make the remainder $e^{it\delta(\xi)}$ non-oscillatory. This would effectively transform each wave-packet into a proper "molecule." For curvelets, see how equation (2.52) depends on the crucial estimate (A.8) about the phase, and how this implies the molecule inequality (2.17).

It is easy to see that $\delta(\xi)$ is zero on the line $\xi = \text{const} \times \mathbf{e}_\mu$, and proportional to $\frac{(\xi \cdot \mathbf{e}_\mu^\perp)^2}{|\xi \cdot \mathbf{e}_\mu|}$ away from it. If $\hat{\varphi}_\mu$ is supported around ξ_μ so that $|\xi_\mu| \sim 2^j$, and the support lies well away from the origin, then $\delta(\xi) \leq \text{const}$ implies that $\xi \cdot \mathbf{e}_\mu^\perp$ be bounded by constant times $2^{j/2}$. This is saying that the width of the support should be at most the square root of the length (in frequency), i.e., $0 \leq \beta \leq 1/2$. In short, *any scaling more anisotropic than the parabolic scaling works.*

In conclusion, only the parabolic scaling, $\beta = 1/2$, allows to formulate a sparsity result like Theorem 1.1 because it meets both requirements of small warping and small dispersion effects.

2.4 Representation of Linear Hyperbolic Systems

We now return to the main theme of this chapter and consider linear initial-value problems of the form

$$\frac{\partial u}{\partial t} + \sum_{k=1}^{m} A_k(x) \frac{\partial u}{\partial x_k} + B(x)u = 0, \qquad u(0,x) = u_0(x), \tag{2.25}$$

where in addition to the properties listed in the introduction, A_k and B together with all their partial derivatives are uniformly bounded for $x \in \mathbb{R}^n$. As explained in Section 2.4.2, we need to make the technical assumption that for every set of real parameters ξ_k, the (real) eigenvalues of the matrix $\sum_k A_k(x)\xi_k$ have constant multiplicity in x and ξ.

Curvelets will provide a concrete "basis" of $L^2(\mathbb{R}^n, \mathbb{C}^m)$ in which the evolution is simple/sparse. We choose to specialize our discussion to $n = 2$ spatial dimensions. The reason is twofold: first, this setting is indeed that in which the exposition of curvelets is the most convenient; and second, this is not a restriction as similar results would hold in arbitrary dimensions.

2.4.1 Main Result

We need to prove

$$|E(t;\mu,\nu;\mu',\nu')| \leq C_{t,N} \cdot \sum_{\nu''} \omega(\mu,\mu'_{\nu''}(t))^{-N}, \tag{2.26}$$

for some constant $C_{t,N} > 0$ growing at most like $C_N e^{K_N t}$ for some $C_N, K_N > 0$. The sum over ν'' indexes the different flows and takes on as many values as there are distinct eigenvalues $\lambda^0_{\nu''}$.

It is instructive to notice that the estimate (2.26) for $t = 0$ is already the strongest of its sort on the off-diagonal decay of the Gram matrix elements for a tight frame of curvelets. For $t > 0$, equation (2.26) states that the strong phase-space localization of every curvelet is preserved by the hyperbolic system, thus yielding a sparse and well organized structure for the curvelet matrix. These warped and displaced curvelets are "curvelet molecules" as introduced in Section 2.2.3 because, as we will show, they obey the estimates (2.12) and (2.13).

The choice of the curvelet family being complex-valued in the above theorem is not essential. $E(t)$ acting on real-valued curvelets would yield *two* molecules per flow (upstream and downstream). Keeping track of this fact in subsequent discussions would be unnecessarily heavy. In the real case, it is clear that the structure and the sparsity of the curvelet matrix can be recovered by expressing each real curvelet as a superposition of two complex curvelets.

The following two sections present results which are for the most part established knowledge in the theory of hyperbolic equations. For example, we borrow some methods and results from geometric optics [57] and most notably from Taylor [84] and Stolk and de Hoop [81]. The goal here is to keep the exposition self-contained and at a reasonable level, and to recast prior results in the framework adopted here, which is sometimes significantly different from that used by the original contributors.

2.4.2 Decoupling into Polarized Components

How to disentangle the vector wavefield into m independent components is perhaps best understood in the special case of constant coefficients, $A_k(x) = A_k$, and with $B(x) = 0$. In this case, applying the 2-dimensional Fourier transform on both sides of (2.25) gives a

system of ordinary differential equations

$$\frac{d\hat{u}}{dt}(t,\xi) + ia(\xi)\hat{u}(t,\xi) = 0, \quad a(\xi) = \sum_k A_k \xi_k.$$

(Note that $a(\xi)$ is a symmetric matrix with real entries.) It follows from our assumptions that one can find m real eigenvalues $\lambda_\nu(\xi)$ and orthonormal eigenvectors $r_\nu(\xi)$, so that

$$a(\xi)r_\nu(\xi) = \lambda_\nu(\xi)r_\nu(\xi).$$

Put $f_\nu(t,\xi) = r_\nu(\xi) \cdot \hat{u}(t,\xi)$. Then our system of equations is of course equivalent to the system of independent scalar equations

$$\frac{df_\nu}{dt}(t,\xi) + i\lambda_\nu(\xi)f_\nu(\xi) = 0,$$

which can then be solved for explicitly;

$$f_\nu(t,\xi) = e^{-it\lambda_\nu(\xi)}f_\nu(0,\xi).$$

Hence, the diagonalization of $a(\xi)$ decouples the original equation (2.25) into m *polarized* components; these can be interpreted as waves going in definite directions, for example "up and down" or "outgoing and incoming" depending on the geometry of the problem. This is the reason why f_ν is also referred to as being a "one-way" wavefield.

The situation is more complicated when $A_k(x)$ is non-uniform since Fourier techniques break down. A useful tool in the variable coefficient setting is the calculus of pseudodifferential operators. An operator T is said to be *pseudodifferential* with symbol σ if it can be represented as

$$Tf(x) = \sigma(x,D)f = \frac{1}{(2\pi)^2}\int_{\mathbb{R}^2} e^{ix\cdot\xi}\sigma(x,\xi)\hat{f}(\xi)\,d\xi, \tag{2.27}$$

with the convention that $D = -i\nabla$. It is of type $(1,0)$ and order m if σ obeys the estimate

$$|\partial_\xi^\alpha \partial_x^\beta \sigma(x,\xi)| \leq C_{\alpha,\beta} \cdot (1+|\xi|)^{m-|\alpha|}$$

for every multi-indices α and β. Unless otherwise stated, all pseudodifferential operators in this thesis are of type $(1,0)$. An operator is said to be smoothing of order $-\infty$, or simply

smoothing if its symbol satisfies the above inequality for every $m < 0$. Observe that this is equivalent to the property that T maps boundedly distributions in the Sobolev space H^{-s} to functions in H^s for every $s > 0$, in addition to a strong localization property of its kernel $G(x, y)$ which says that for each $N > 0$, there is a constant $C_N > 0$ such that G obeys

$$|G(x, y)| \leq C_N \cdot (1 + |x - y|)^{-N} \tag{2.28}$$

as in [79](Chapter 6).

Now set

$$a(x, D) = \sum_{k=1}^{m} A_k(x) D_k - iB(x),$$

and its principal part

$$a^0(x, D) = \sum_{k=1}^{m} A_k(x) D_k,$$

so that equation (2.25) becomes $\partial_t u + ia(x, D)u = 0$. The matrices $a(x, \xi)$ (resp. $a^0(x, \xi)$) are called the symbol of the operator $a(x, D)$ (resp. $a^0(x, D)$). Note that $a^0(x, \xi)$ is homogeneous of degree one in ξ; a^0 also goes by the name of *dispersion matrix.*

It follows from the symmetry of A_k and B that for every set of real parameters $\xi_1, \ldots, \xi_m$, the matrix $a^0(x, \xi) = \sum_k A_k(x)\xi_k$ is also symmetric and thus admits real eigenvalues $\lambda^0_\nu(x, \xi)$ and an orthonormal basis of eigenvectors $r^0_\nu(x, \xi)$,

$$a^0(x, \xi) r^0_\nu(x, \xi) = \lambda^0_\nu(x, \xi) r^0_\nu(x, \xi). \tag{2.29}$$

The eigenvalues being real and the set of eigenvectors complete is a *hyperbolicity* condition and ensures that equation (2.25) will admit wave-like solutions. We assume throughout this thesis that the multiplicity of each $\lambda^0_\nu(x, \xi)$ is constant in x and ξ.

By analogy with the special case of constant coefficients, a first impulse may be to introduce the components $r^0_\nu(x, D) \cdot u$, where $r^0_\nu(x, D)$ is the operator associated to the eigenvector $r^0_\nu(x, \xi)$ by the standard rule (2.27). In particular this is how we defined hypercurvelets from curvelets in Section 1.2.2. Unfortunately, this does not perfectly decouple the system into m polarized modes—it only approximately decouples. Instead, we would achieve perfect decoupling if we could solve the eigenvalue problem

$$a(x, D) r_\nu(x, D) = r_\nu(x, D) \lambda_\nu(x, D). \tag{2.30}$$

Here, each $\Lambda_\nu = \lambda_\nu(x, D)$ is a scalar operator and $R_\nu = r_\nu(x, D)$ is an m-by-1 vector of operators. Equation (2.30) must be understood in the sense of composition of operators. Now let f_ν be the polarized components obeying the scalar equation

$$\frac{\partial f_\nu}{\partial t} + i\Lambda_\nu f_\nu = 0, \tag{2.31}$$

with initial condition $f_\nu(0, x)$ and consider the superposition

$$u = \sum_\nu u_\nu, \quad u_\nu = R_\nu f_\nu.$$

Then u is a solution to our initial-value problem (2.25). (We will make this rigorous later, and detail the dependence between the initial values u_0 and the $f_\nu(0, \cdot)$.)

The following result shows how in some cases, (2.30) can be solved up to a smoothing remainder of order $-\infty$. When all the eigenvalues $\lambda^0_\nu(x, \xi)$ are simple, the exact diagonalization is, in fact, possible. The situation is more complicated when some of the eigenvalues are degenerate; further decoupling within the eigenspaces is in general not possible. This complication does not compromise, however, any of our results.

The theorem is due to Taylor [84], Stolk and de Hoop [81].

Lemma 2.3. *Suppose our hyperbolic system satisfies all the assumptions stated below (2.25). Then there exists an m-by-m block-diagonal matrix of operators Λ and two m-by-m matrices of operators R and S such that*

$$a(x, D)R = R\Lambda + S,$$

where Λ, R and S are componentwise pseudodifferential with Λ of order one, R of order zero, and S of order $-\infty$. Each block of Λ corresponds to a distinct eigenvalue λ^0_ν whose size equals the multiplicity of that eigenvalue. The principal symbol of Λ is diagonal with the eigenvalues $\lambda^0_\nu(x, \xi)$ as entries.

Let us provide an alternative, easier proof of the Taylor-Stolk-de Hoop lemma.

Proof. We already argued (2.30) is not just the eigenvalue problem for the symbol $a(x, \xi)$ for the composition of two operators does not reduce to a multiplication of their respective symbols. Instead, it is common practice [41] to define the *twisted product* of two symbols σ

and τ as

$$(\sigma \sharp \tau)(x, D) = \sigma(x, D)\tau(x, D),$$

so that (2.30) becomes the symbol equation $a \sharp r_\nu = r_\nu \sharp \lambda_\nu$. Note that $D = -i\nabla$. The explicit formula for the twisted product is, in multi-index notation[2],

$$\sigma \sharp \tau = \sum_{|\alpha| \geq 0} \frac{1}{\alpha!} \partial_\xi^\alpha \sigma D_x^\alpha \tau.$$

We can see that $\sigma \sharp \tau$ is the product $\sigma\tau$ up to terms that are at least one order lower (because of the differentiations in ξ).

Recall the decomposition of the symbol $a(x, \xi)$ into a principal part $a^0(x, \xi) = \sum_k A_k(x)\xi_k$, homogeneous of degree one in ξ, and a remainder $B(x)$ homogeneous of order zero. It follows that the eigenvalues $\lambda_\nu^0(x, \xi)$ of $a^0(x, \xi)$ are homogeneous of degree one, and the corresponding eigenvectors $r_\nu^0(x, \xi)$ may be selected as homogeneous of degree zero (and orthonormal). Up to terms of lower order in ξ, the original problem (2.30) therefore reduces to the eigenvalue problem $a^0(x, \xi)r_\nu^0(x, \xi) = r_\nu^0(x, \xi)\lambda_\nu^0(x, \xi)$ for the symbol a^0. It is then natural to look for a solution r_ν, λ_ν of (2.30) as a perturbation of r_ν^0, λ_ν^0 by lower-order terms.

Consider first the case in which each eigenvalue λ_ν^0 is simple and define the expansions

$$r_\nu \sim r_\nu^0 + r_\nu^1 + r_\nu^2 + \ldots, \qquad \lambda_\nu \sim \lambda_\nu^0 + \lambda_\nu^1 + \lambda_\nu^2 + \ldots$$

so that r_ν^n is of order $-n$ in ξ and λ_ν^n of order $-n + 1$ i.e.,

$$|\partial_\xi^\alpha \partial_x^\beta r_\nu^n(x, \xi)| \leq C_{\alpha,\beta}(1 + |\xi|)^{-n-|\alpha|},$$

and similarly for λ_ν^n. We plug these expansions in the twisted product, or equivalently in (2.30), and isolate terms of identical degree.

The contribution at the leading order is, of course, $a^0 r_\nu^0 = \lambda_\nu^0 r_\nu^0$ and the remainder is of the form $a \sharp r_\nu^0 - r_\nu^0 \sharp \lambda_\nu^0$; put e_ν^0 as its principal symbol. The zero-order equation reads

$$(a^0 - \lambda_\nu^0 I)r_\nu^1 = -e_\nu^0 + r_\nu^0 \lambda_\nu^1, \tag{2.32}$$

[2]All the pseudodifferential operators considered in this paper are of type $(1, 0)$ therefore all such polyhomogeneous expansions are valid.

which admits a solution if and only if the right hand side has a zero component in the eigenspace spanned by r_ν^0. This is possible if λ_ν^1 is selected so that

$$-e_\nu^0 + r_\nu^0 \lambda_\nu^1 \;\perp\; r_\nu^0 \quad \Leftrightarrow \quad \lambda_\nu^1 = r_\nu^0 \cdot e_\nu^0.$$

It follows that equation (2.32) admits the family of solutions

$$r_\nu^1 = (a^0 - \lambda_\nu^0 I)^{-1}(-e_\nu^0 + r_\nu^0 \lambda_\nu^1) + f_1 r_\nu^0,$$

where f_1 is actually a scalar function of x and ξ, and homogeneous of degree -1 in ξ. Our proof does not exploit this degree of freedom.

It is clear that one can successively determine all the λ_ν^n's and r_ν^n's in a similar fashion. Let e_ν^n be the principal symbol of $a \sharp (r_\nu^0 + \ldots + r_\nu^n) - (r_\nu^0 + \ldots + r_\nu^n) \sharp (\lambda_\nu^0 + \ldots + \lambda_\nu^n)$, then the equation at the order $-n$ is

$$(a^0 - \lambda_\nu^0 I) r_\nu^{n+1} = -e_\nu^n + r_\nu^0 \lambda_\nu^{n+1},$$

and is solved exactly like (2.32).

Suitable cutoffs of the low frequencies guarantee convergence of the series for r_ν and λ_ν. As is standard in the theory of pseudodifferential operators [75, 85], one selects a sequence of C^∞ cut-off functions $\chi_n(\xi) = \chi(\epsilon^n \xi)$ for some χ vanishing inside a compact neighborhood of the origin, and identically equal to one outside a larger neighborhood. Then ϵ is taken small enough so that

$$r_\nu(x,\xi) = \sum_{n=0}^{\infty} r_\nu^n(x,\xi)\chi_n(\xi), \qquad \lambda_\nu(x,\xi) = \sum_{n=0}^{\infty} \lambda_\nu^n(x,\xi)\chi_n(\xi)$$

are converging expansions in the topology of C^∞. As a result, the remainder $s_\nu = a \sharp r_\nu - r_\nu \sharp \lambda_\nu$ also converges to a valid symbol which, by construction, is of order $-\infty$, i.e., obeys

$$|\partial_\xi^\alpha \partial_x^\beta s_\nu(x,\xi)| \leq C_{\alpha,\beta,N} \cdot (1 + |\xi|)^{-N-|\alpha|}$$

for every $N > 0$. The lemma is proved in the case when all eigenvalues of the principal symbol are simple.

Consider now the case of a multiple eigenvalue λ^0, say. Suppose the corresponding

eigenspace is of dimension p and spanned by $r_1^0, \dots, r_p^0$. The reasoning for simple eigenvalues does not apply because the p solvability conditions are too many for purely diagonal lower-order corrections. Instead, the block corresponding to λ^0 is now perturbed as

$$\begin{pmatrix} \lambda^0 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \lambda_0 \end{pmatrix} + \begin{pmatrix} \lambda_{11}^1 & \cdots & \lambda_{1p}^1 \\ \vdots & \ddots & \vdots \\ \lambda_{p1}^1 & \cdots & \lambda_{pp}^1 \end{pmatrix} + \begin{pmatrix} \lambda_{11}^2 & \cdots & \lambda_{1p}^2 \\ \vdots & \ddots & \vdots \\ \lambda_{p1}^2 & \cdots & \lambda_{pp}^2 \end{pmatrix} + \dots$$

where each λ_{ij}^n is homogeneous of degree $-n+1$ in ξ. At the leading order, The p equations relative to λ_0 are

$$(a^0 - \lambda^0 I) r_j^1 = -e_j^0 + \sum_{i=1}^{p} r_i^0 \lambda_{ij}^1, \tag{2.33}$$

where e_j^0 is the principal symbol of $a \sharp r_j^0 - r_j^0 \sharp \lambda^0$. Solvability requires that the projection of the right-hand side on each of the r_i^0, $i = 1, \dots, p$ vanishes. This unambiguously determines all the components of the p-by-p block λ^1 as

$$\lambda_{ij}^1 = r_i^0 \cdot e_j^0.$$

All blocks relative to other eigenvalues are solved for in a similar way, yielding a block-diagonal structure for the zeroth order correction λ^1. Each block should have dimension equal to the multiplicity of the corresponding eigenvalue in order to meet the solvability requirements.

The perturbed eigenvectors $r_1^1, \dots, r_p^1$ are determined as previously once the λ_{ij}^1 are known. The same reasoning applies at all orders and thereby determines Λ and R. Convergence issues are addressed using cutoff windows just as before.

□

The above construction indeed provides efficient decoupling of the original problem (2.25) into polarized modes. The following lemma is a straightforward consequence of Lemma 2.3 although we have not been able to find it in the literature. See [81] for related results.

Lemma 2.4. *In the setting of Lemma 2.3, the solution operator $E(t)$ for (2.25) may be*

decomposed for all times $t > 0$ as

$$E(t) = Re^{-it\Lambda}L + \tilde{S}(t),$$

where the matrices of operators Λ and R are defined in Lemma 2.3 and $\tilde{S}(t)$ is (another) matrix of smoothing operators of order $-\infty$. In addition,

1. *L is an approximate inverse of R, i.e., $RL = I$ and $LR = I$ (mod smoothing).*

2. *L is a pseudodifferential of order zero (componentwise).*

Observe that $e^{-it\Lambda}$ inherits the block structure from Λ, and is diagonal in the case where all the eigenvalues λ_ν^0 are simple.

Proof. Begin by observing that $R = r(x, D)$—as an operator acting on $L^2(\mathbb{R}^2, \mathbb{C}^m)$—is invertible modulo a smoothing additive term. This means that one can construct a parametrix L so that $LR = I$ and $RL = I$ with both equations holding modulo a smoothing operator. To see why this is true, note that the matrix $r(x, \xi)$ is a lower-order perturbation from the unitary matrix $r^0(x, \xi)$ of eigenvectors of the principal symbol $a^0(x, \xi)$. The inverse of $r^0(x, \xi)$ is explicitly given by $\ell^0(x, \xi) = r^0(x, \xi)^*$. The symbol of L can now be built as an expansion $\ell^0 + \ell^1 + \ldots$, where each $\ell^n(x, \xi)$ is homogeneous of degree $-n$ in ξ and chosen to suppress the $O(|\xi|^{-j})$ contribution in $RL - I$ as well as in $LR - I$. This construction implies that L is pseudodifferential of order zero (componentwise). All of this is routine and detailed in [41](page 117).

In the sequel, S, S_1 and S_2 will denote a generic smoothing operator whose value may change from line to line. The composition of a pseudodifferential operator and a smoothing operator is obviously still smoothing. Set $f = Lu$ and let $A = a(x, D)$, so that $\partial_t u = -iAu$. On the one hand, $u = Rf - Su$ and

$$\partial_t u = R\partial_t f - S\partial_t u = R\partial_t f - SAu. \tag{2.34}$$

On the other hand, Lemma 2.3 gives

$$Au = ARf - ASu = R\Lambda f + S_1 f + S_2 u = R\Lambda f + Su \tag{2.35}$$

Comparing (2.34) and (2.35), and applying L gives

$$\partial_t f = -i\Lambda f + Su. \tag{2.36}$$

This can be solved by Duhamel's formula,

$$f(t) = e^{-it\Lambda} f(0) + \int_0^t e^{-i(t-\tau)\Lambda} Su(\tau)\, d\tau. \tag{2.37}$$

We now argue that the integral term is, indeed, a smoothing operator applied to the initial value u_0.

- First, the evolution operator $E(t) = e^{-itA}$ has a kernel $K(t, x, y)$ supported inside a neighborhood of the diagonal $y = x$ and for each $s \geq 0$, is well known to map $H^s(\mathbb{R}^2, \mathbb{C}^m)$ boundedly onto itself [57]. Therefore, $SE(\tau)$ maps H^{-s} to H^s boundedly for every $s > 0$ and has a well localized kernel in the sense of (2.28). This implies that $SE(\tau)$ is a smoothing operator.

- Second, Section 2.4.3 shows that $e^{-it\Lambda}$ is, for small t, a FIO of type $(1, 0)$ and order zero, modulo a smoothing remainder. The composition of a FIO and a smoothing operator is smoothing. For larger t, think about $e^{-it\Lambda}$ as the product $(e^{-i\frac{t}{n}\Lambda})^n$ for appropriately large n.

- And third, the integral extends over a finite interval $[0, t]$ and may be thought as an average of smoothing operator—hence smoothing.

In short, $f(t) = e^{-it\Lambda} f(0) + Su_0$. Applying R on both sides of (2.36) finally gives

$$u = Re^{-it\Lambda} Lu_0 + S_1 u_0 + S_2 u = (Re^{-it\Lambda} L + S)u_0$$

which is what we set out to establish.

It remains to see that the evolution operator $e^{-it\Lambda}$ for the polarized components has the same block-diagonal structure as Λ itself. This is gleaned from equation (2.36): evolution equations for two components f_{ν_1}, f_{ν_2} (corresponding to distinct eigenvalues) are completely decoupled. □

2.4.3 The Fourier Integral Operator Parametrix

Lemma 2.4 explained how to turn the evolution operator $E(t)$ into the block-diagonal representation $e^{-it\Lambda}$. In this section, we describe how each of these blocks can be approximated by a Fourier integral operator. The ideas here are standard and our exposition is essentially taken from [43] and [75]. The original construction is due to Lax [57].

Let us first assume that all eigenvalues of the principal symbol $a^0(x,\xi)$ are simple. This is the situation where the matrix of operators Λ (Lemma 2.3) is diagonal with elements Λ_ν. Put $E_\nu(t) = e^{-it\Lambda_\nu}$, the (scalar) evolution operator relative to the νth polarized mode. We seek a parametrix $F_\nu(t)$ such that $S_\nu(t) = E_\nu(t) - F_\nu(t)$ is smoothing of order $-\infty$.

Formally,

$$f(t,x) = \int E_\nu(t)(e^{ix\cdot\xi})\ \widehat{f_0}(\xi)\ d\xi.$$

Our objective is to build a high-frequency asymptotic expansion for $E_\nu(t)(e^{ix\cdot\xi})$ of the form

$$e^{i\Phi_\nu(t,x,\xi)}\sigma_\nu(t,x,\xi), \tag{2.38}$$

where $\sigma_\nu \sim \sigma_\nu^0 + \sigma_\nu^1 + \dots$ with σ_ν^n homogeneous of degree $-n$ in ξ, and Φ homogeneous of degree one in ξ.

As is classical in asymptotic analysis, we proceed by applying $M_\nu = \partial_t + i\Lambda_\nu$ to the expansion (2.38) and successively equate all the coefficients of the negative powers of $|\xi|$ to zero, hence mimicking the relation $M_\nu E_\nu(t)(e^{ix\cdot\xi}) = 0$ which holds by definition. For obvious reasons, we also impose that (2.38) evaluated at $t = 0$ be $e^{ix\cdot\xi}$. Note that, in accordance to Lemma 2.3, Λ_ν is taken as a polyhomogeneous expansion $\sum_{j\geq 0}\lambda_\nu^j(x,D)$, where each symbol $\lambda_\nu^j(x,\xi)$ is homogeneous of degree $-j+1$ in ξ.

After elementary manipulations, one finds that the phases must satisfy the standard Hamilton-Jacobi equations

$$\frac{\partial\Phi_\nu}{\partial t} + \lambda_\nu^0(x,\nabla_x\Phi_\nu) = 0, \tag{2.39}$$

with $\Phi_\nu(0,x,\xi) = x\cdot\xi$. The amplitudes σ_ν^n are successively determined as the solutions of transport equations along each Hamiltonian vector field,

$$\frac{\partial\sigma_\nu^n}{\partial t} + \nabla_\xi\lambda_\nu^0(x,\nabla_x\Phi_\nu)\cdot\nabla_x\sigma_\nu^n = P_n(\sigma_\nu^0,\dots,\sigma_\nu^n), \tag{2.40}$$

where P_n is a known differential operator applied to $\sigma_\nu^0, \ldots, \sigma_\nu^n$.

In the case where some eigenvalue λ_ν^0 has multiplicity $p > 1$, the construction of a FIO parametrix goes the same way, except that Λ_ν denotes the p-by-p block corresponding to λ_ν^0 in the matrix Λ of Lemma 2.3. Also each σ_ν^n is now a p-by-p matrix of amplitudes.

It is important to notice that Φ_ν may be defined only for small times, because it would become multivalued when rays originating from the same point x_0 cross again later. This typically happens at cusp points, when caustics start developing. We refer the interested reader to [43, 89].

We skipped a lot of justifications in the above exposition, in particular on convergence issues, but these technicalities are standard and detailed in some very good monographs. The following result summarizes all that we shall need.

Lemma 2.5. *Define t^* as half the infimum time for which a solution to (2.39) ceases to exist, uniformly in ν and ξ. In the setting of Lemma 2.3, denote by Λ_ν a block of Λ and $E_\nu(t) = e^{-it\Lambda_\nu}$. Then for every $0 < t \leq t^*$, there exists a parametrix $F_\nu(t)$ for the evolution problem $\partial_t f + i\Lambda_\nu f = 0$ which takes the form of a Fourier integral operator,*

$$F_\nu(t) f_0(x) = \int e^{i\Phi(t,x,\xi)} \sigma_\nu(t,x,\xi) \hat{f}_0(\xi)\, d\xi.$$

For each $t \leq t^*$, the phase function Φ_ν is positive-homogeneous of degree one in ξ and smooth in x and ξ; the amplitude σ_ν is a symbol of type $(1,0)$ and order zero. The remainder $S_\nu(t) = E_\nu(t) - F_\nu(t)$ is a smoothing operator of order $-\infty$.

Proof. The proof is for the most part presented in [75](Pages 120 and below). See also [36, 43, 85]. □

2.4.4 Sparsity of Smoothing Terms

The specialist will immediately recognize that a smoothing operator of order $-\infty$ is very sparse in a curvelet frame. This is the content of the following lemma.

Lemma 2.6. *The curvelet entries of a smoothing operator S obey the following estimate: for each $N > 0$, there is a constant C_N such that*

$$|\langle \varphi_\mu, S\varphi_{\mu'} \rangle| \leq C_N \cdot 2^{-|j+j'|N} (1 + |x_\mu - x_{\mu'}|)^{-N}. \tag{2.41}$$

Note that (2.41) is a stronger estimate than that of Theorem 1.1. Indeed, our lemma implies that

$$|\langle \varphi_\mu, S\varphi_{\mu'}\rangle| \le C_N \cdot \omega(\mu, \mu'_\nu(t))^{-N}$$

which is valid for each $N > 0$ and regardless of the value of ν.

Proof. We know that S maps H^{-s} to H^s for arbitrary large s, so does its adjoint S^*. As a result,

$$\begin{aligned} |\langle \varphi_\mu, S\varphi_{\mu'}\rangle| &\le |\langle S^*\varphi_\mu, \varphi_{\mu'}\rangle|^{1/2} |\langle \varphi_\mu, S\varphi_{\mu'}\rangle|^{1/2} \\ &\le \|S^*\varphi_\mu\|_{H^s}^{1/2} \|\varphi_{\mu'}\|_{H^{-s}}^{1/2} \|\varphi_\mu\|_{H^{-s}}^{1/2} \|S\varphi_{\mu'}\|_{H^s}^{1/2} \\ &\le C \cdot \|\varphi_{\mu'}\|_{H^{-s}} \|\varphi_\mu\|_{H^{-s}} \le C \cdot 2^{-(j+j')s}. \end{aligned}$$

Next, recall that curvelets have an essential spatial support of size at most $O(1) \times O(1)$. (Coarse scale curvelets have support size about $O(1) \times O(1)$ and the size decreases at increasingly finer scales.) The action of S is local on this range of distances, so that

$$|\langle \varphi_\mu, S\varphi_{\mu'}\rangle| \le C_N \cdot (1 + |x_\mu - x_{\mu'}|)^{-N}.$$

for arbitrary large $N > 0$. These two bounds can be combined to conclude that the matrix elements of S are negligible in the sense defined above. □

2.4.5 Proof of Theorem 1.1

Let us first show how the first assertion on the near-exponential decay of the curvelet matrix elements follows immediately from the second one, equation (1.19). Let a be either a row or a column of the curvelet matrix and let $|a|_{(n)}$ be the nth largest entry of the sequence $|a|$. We have

$$n^{1/p} \cdot |a|_{(n)} \le \|a\|_{\ell_p}^{1/p}$$

and, therefore, it is sufficient to prove that the matrix E has rows and columns bounded in ℓ_p for every $p > 0$. Consider the columns. We need to establish

$$\sup_{\mu',\nu'} \sum_{\mu,\nu} |E(t; \mu, \nu; \mu', \nu')|^p \le C_{t,p}, \tag{2.42}$$

for some constant $C_{t,p} > 0$ growing at most like $C_p e^{K_p t}$ for some $C_p, K_p > 0$.

The sum over ν and the sup over ν' do not come in the way since these subscripts take on a finite number of values. The fine decoupling between the m one-way components, crucial for equation (1.19), does not play any role here.

Let us now show that there exists N so that

$$\sum_{\mu} \omega(\mu, \mu')^{-Np} \leq C_{N,p},$$

uniformly in μ'. We can use the bound (A.2) with Np in place of N for the sum over k and ℓ. This gives

$$\sum_{\mu} \omega(\mu, \mu')^{-Np} \leq C_{N,p} \cdot \sum_{j \geq 0} 2^{-|j-j'|Np} \cdot 2^{2|j-j'|},$$

which is bounded by a constant depending on N and p provided again that $Np \geq 2$.

Hence we proved the property for the columns. The same holds for the rows because the same conclusion is true for the adjoint $E(t)^*$; indeed, the adjoint solves the backward initial-value problem for the adjoint equation $u_t = A^* u$, and A^* satisfies the same hyperbolicity conditions as A. We can therefore interchange the role of the two curvelets and obtain

$$\sup_{\mu', \nu'} \sum_{\mu, \nu} |E(t; \mu', \nu'; \mu, \nu)|^p \leq C_{t,p}.$$

Note that the classical interpolation inequality shows that $E(t)$ is a bounded operator from ℓ_p to ℓ_p for every $0 < p \leq \infty$.

We now turn to (1.19). Let us assume first that all eigenvalues λ_ν^0 of the principal symbol a^0 are simple. According to Lemmas 2.4 and 2.5, each matrix element $E(t; \nu; \nu') = \mathbf{e}_\nu \cdot E(t) \mathbf{e}_{\nu'}$ of $E(t)$ can for fixed (possibly large) time $t > 0$ be written as

$$E(t; \nu; \nu') = \sum_{\nu''=1}^{m} R_{\nu,\nu''} (e^{-i\frac{t}{n}\Lambda_{\nu''}})^n L_{\nu'',\nu'} + S_{\nu,\nu'}(t). \tag{2.43}$$

We have taken n large enough—proportional to t—so that $e^{-i\frac{t}{n}\Lambda_\nu}$ is a Fourier integral operator (mod smoothing) for every ν. Each $R_{\nu,\nu'}$ and $L_{\nu,\nu'}$ is pseudodifferential of order zero and $S_{\nu,\nu'}(t)$ is smoothing.

Thanks to Lemma 2.6, we only need to prove the claim for the first term of (2.43) which follows from Theorem 2.1 in Section 2.5 about the sparsity of FIOs in a curvelet tight frame.

As is well known, the ray dynamics is equivalently expressed in terms of Hamiltonian flows

$$\begin{cases} \dot{x}(t) &= \nabla_\xi \lambda^0(x(t),\xi(t)), \qquad & x(0)=x_0, \\ \dot{\xi}(t) &= -\nabla_x \lambda^0(x(t),\xi(t)), \qquad & \xi(0)=\xi_0, \end{cases} \tag{2.44}$$

or in terms of canonical transformations generated by the phase functions Φ_ν,

$$\begin{cases} x_0 &= \nabla_\xi \Phi(t,x(t),\xi_0), \\ \xi(t) &= \nabla_x \Phi(t,x(t),\xi_0), \end{cases} \tag{2.45}$$

provided $\Phi(t,x,\xi)$ satisfies the Hamilton-Jacobi equation $\frac{\partial \Phi}{\partial t} + \lambda^0(x,\nabla_x\Phi) = 0$ with initial condition $\Phi(0,x,\xi_0) = x\cdot\xi_0$. We obviously need this property to ensure that the geometry of FIOs is the same as that of hyperbolic equations.

Pseudodifferential operators are a special instance of Fourier integral operators so the theorem equally applies to them. For $E(t;\mu,\nu;\mu',\nu') = \langle \varphi_\mu, E(t;\nu;\nu')\varphi'_\mu\rangle$ we get

$$|E(t;\mu,\nu;\mu',\nu')| \leq C_N \sum_{\nu''=1}^{m} \sum_{\mu_0} \cdots \sum_{\mu_n} \omega(\mu,\mu_0)^{-N} \omega(\mu_0,\mu_{1\nu''}(\frac{t}{n}))^{-N} \cdots \\ \omega(\mu_{n-1},\mu_{n\nu''}(\frac{t}{n}))^{-N} \omega(\mu_n,\mu')^{-N},$$

for all $N>0$. Inequality (2.26) then follows from repeated applications of properties 3 and 4 of the distance ω, see proposition 2.1. The power growth in t of the overall multiplicative constant comes from the number of intermediate sums over $\mu_0,\ldots,\mu_n$. There are $n+1\sim t$ such sums and they each introduce the same multiplicative constant C_N.

The reasoning is the same when at least some eigenvalues λ^0_ν are degenerate. The subscript ν'' now denotes the flows i.e., the eigenvalues $\lambda^0_{\nu''}$ *not* counting their multiplicity. Each $R_{\nu,\nu''}$ is a row vector, $e^{-i\frac{t}{n}\Lambda_{\nu''}}$ a matrix and $L_{\nu'',\nu'}$ a column vector. The FIO parametrix for $e^{-i\frac{t}{n}\Lambda_{\nu''}}$ was constructed in such a way that only one flow $h_{\nu''}$ appears in the majoration of its curvelet elements (componentwise). There is no intermediate sum over $\nu_0,\ldots,\nu_n$ and this is the whole point of decoupling the polarized components *before* constructing the FIO parametrix.

2.4.6 Relation to Hypercurvelets

In Section 1.2.2 we introduced hypercurvelets as "polarized" curvelets which would not split into m molecules along the m different flows. In light of Section 2.4.2, it is interesting to reformulate our main result (2.26) in terms of hypercurvelets. We recall that

$$\boldsymbol{\varphi}^{(0)}_{\mu\nu}(x) = r^0_\nu(x,D)\varphi_\mu(x) = \frac{1}{(2\pi)^2}\int e^{ix\cdot\xi} r^0_\nu(x,\xi)\hat{\varphi}_\mu(\xi)\,d\xi.$$

Corollary 2.1. *Define* $E^{(0)}(t;\mu,\nu;\mu',\nu') = \langle \boldsymbol{\varphi}^{(0)}_{\mu\nu}, E(t)\boldsymbol{\varphi}^{(0)}_{\mu'\nu'}\rangle$. *Then under the same assumptions as those of Theorem 1.1 we have for all* $N > 0$

$$|E^{(0)}(t;\mu,\nu;\mu',\nu')| \le C_{tN}\cdot[\,\omega(\mu,\mu'_{\nu'}(t))^{-N} + 2^{-j'}\sum_{\nu''\neq\nu'}\omega(\mu,\mu'_{\nu''}(t))^{-N}\,]. \tag{2.46}$$

The main contribution to the right-hand side is due to the ν'th flow. All other flows are weighted by the small factor $2^{-j'}$ (which is about equal to $|\xi|^{-1}$ on the support of $\hat{\varphi}_{\mu'}$). In other words, there might be some "cross talk" between the various components corresponding to the different flows but it is at most smoothing of order -1, hence small at small scales.

Proof. Equation (2.46) follows from Theorem 1.1 and the fact that the adjoint of the matrix operator R^0 whose columns are the $R^0_\nu = r^0_\nu(x,D)$ is an approximate left inverse for R^0—up to an error smoothing of order -1. Indeed, by the standard rules for composition and computation of the adjoint of pseudodifferential operators,

$$\begin{aligned}(R^0_\nu)^* R^0_{\nu'} &= ((r^0_\nu)^* \sharp\, r^0_{\nu'})(x,D)\\ &= ((r^0_\nu)^* r^0_{\nu'})(x,D) + \text{ order } -1\\ &= \delta_{\nu\nu'} I + \text{ order } -1.\end{aligned}$$

We have used the fact that the dispersion matrix $a^0(x,\xi)$ is assumed to be symmetric, hence admits an orthobasis of eigenvectors $r^0_\nu(x,\xi)$. We then conclude from Theorem 2.1 applied to pseudodifferential operators of order -1. □

Alternatively, we could have defined hypercurvelets as

$$\boldsymbol{\varphi}_{\mu\nu}^{(\infty)} = r_\nu(x, D)\varphi_\mu(x) = \frac{1}{(2\pi)^2}\int e^{ix\cdot\xi} r_\nu(x,\xi)\hat{\varphi}_\mu(\xi)\, d\xi.$$

This would have given the same result.[3] The reason why we did not use hypercurvelets in the preceding sections is that they do not necessarily constitute a suitable practical basis to decompose wavefields onto. We do not even know if they always constitute a frame. Digital implementation would also seem less obvious.

2.5 Representation of FIOs

The purpose of this section is to show that Fourier integral operators admit a sparse and well organized structure in a curvelet frame. The main result, Theorem 2.1, is a key step in completing the discussion of the previous section. (Observe that by construction, the FIOs encountered in the previous section satisfy all the assumptions stated in Section 2.1 right below (2.1).) As in the previous section, we will restrict the discussion to $x \in \mathbb{R}^2$ which is no loss of generality, see Section 2.6.

2.5.1 Main Results

In the introduction section, we detailed a notion of Hamiltonian correspondence for hyperbolic equations. This correspondence also exists for FIOs and is "encoded" in the phase function Φ of the FIO. It is called the canonical transformation associated to Φ, and is defined as the mapping $(x,\xi) \to (y,\eta)$ of phase-space

$$x = \nabla_\xi \Phi(y,\xi), \qquad \eta = \nabla_x \Phi(y,\xi). \tag{2.47}$$

As suggested in Section 2.4.5, this formulation is equivalent to that involving trajectories along the bicharacteristic flow as in equation (1.10), provided the phase function solves an appropriate Hamilton-Jacobi equation.

This canonical transformation induces a mapping of curvelet subscripts, denoted by $\mu' = h(\mu)$. It is defined via the closest point $(x_{\mu'}, \xi_{\mu'})$, on the curvelet lattice, to the image of (x_μ, ξ_μ) by the canonical transformation. We can already remark that mistaking a point

[3] We can only conjecture that the decoupling should be better if we use the improved $\boldsymbol{\varphi}_{\mu\nu}^{(\infty)}$.

$(x_{\mu'}, \xi_{\mu'})$ for one of its neighbors will not compromise the following result, only increase the value of the constant C_N in front of the estimate.

The main result for this section reads as follows.

Theorem 2.1. *Let T be a Fourier integral operator of order m acting on functions of $\mathbb{R}^2$, with the assumptions stated above, and $T(\mu;\mu')$ denote its matrix elements in the complex curvelet tight frame. Then with h the curvelet index mapping and ω the distance defined in (2.19), the elements $T(\mu;\mu')$ obey for each $N > 0$*

$$|T(\mu;\mu')| \leq C_N \cdot 2^{mj'} \omega(\mu, h(\mu'))^{-N},$$

for some $C_N > 0$. Moreover, for every $0 < p \leq \infty$, $(T(\mu,\mu'))$ is bounded from ℓ^p to ℓ^p.

The interpretation of Theorem 2.1 is in strong analogy with that of Theorem 1.1. Namely, a FIO has the property of transporting and warping a curvelet into another curvelet-like molecule. (Again, the choice of using complex-valued curvelets is not essential, as a real curvelet would be mapped onto two molecules.)

The proof of Theorem 2.1 relies on the factorization of T on the space-frequency support of φ_μ as a nice pseudolocal operator $T_{1,\mu}$ followed by a smooth change of variables, or warping $T_{2,\mu}$. This decomposition goes as follows.

Let φ_μ be a fixed curvelet centered around the lattice point (x_μ, ξ_μ) in phase-space. The phase of our FIO can be decomposed as

$$\Phi(x,\xi) = \Phi_\xi(x,\xi_\mu) \cdot \xi + \delta(x,\xi), \qquad \phi_\mu(x) = \Phi_\xi(x,\xi_\mu). \tag{2.48}$$

In effect, the above decomposition "linearizes" the frequency variable and is classical, see [68, 79]. With these notations, we may rewrite the action of T on a curvelet φ_μ as

$$(T\varphi_\mu)(x) = \int e^{i\phi_\mu(x)\cdot\xi} e^{i\delta(x,\xi)} \sigma(x,\xi) \hat{\varphi}_\mu(\xi)\, d\xi. \tag{2.49}$$

Now for a fixed value of the parameter μ, we introduce the decomposition

$$T = T_{2,\mu} T_{1,\mu},$$

where

$$(T_{1,\mu}f)(x) = \int e^{ix\cdot\xi} b_\mu(x,\xi)\hat{f}(\xi)\,d\xi, \quad (T_{2,\mu}f)(x) = f(\phi_\mu(x)), \tag{2.50}$$

with $b_\mu(x,\xi) = e^{i\delta(\phi_\mu^{-1}(x),\xi)}\sigma(\phi_\mu^{-1}(x),\xi))$. This decomposition allows the separate study of the nonlinearities in frequency ξ and space x in the phase function Φ. The point is that both $T_{1,\mu}$ and $T_{2,\mu}$ are sparse in a curvelet tight frame—only for very different reasons.

Theorem 2.2. *Let $(\varphi_\mu)_\mu$ be a tight frame of curvelets compactly supported in frequency. For each μ, $T_{1,\mu}$ maps φ_μ into a curvelet molecule m_μ with arbitrary regularity R, uniformly over μ in the sense that the constants in estimates (2.12) and (2.13) do not depend on μ.*

As we shall see, the proof of Theorem 2.2, presented in Section 2.5.2, relies on the property of compact support in frequency of the φ_μ. In contrast the corresponding result for the operators $T_{2,\mu}$ which we present next, is extraordinarily simplified if one uses curvelets compactly supported in space. Although well localized in space, the tight frame introduced in Section 2.2 does not meet this requirement. In order to circumvent this technical difficulty, we introduce compactly supported *curvelet atoms* in Section 2.5.3. They are built on the model of atomic decompositions, standard in approximation theory [42].

Theorem 2.3. *Let $(\rho_\mu)_\mu$ be a family of complex-valued curvelet atoms, compactly supported in space, with regularity R. Denote by h the canonical index correspondence associated to Φ, as defined above. For each μ, $T_{2,\mu}$ maps ρ_μ into a molecule $m_{h(\mu)}$ of the same regularity R, uniformly over μ.*

The latter theorem says that the "warped" atom $\rho_\mu \circ \phi_\mu$ is still an atom, only its scale, orientation, and location may have been changed. That a smooth warping preserves the sparsity of curvelet expansions is a result of independent interest.

The remaining three sections are devoted to the proofs of Theorems 2.2, 2.3 and 2.1. The dependence of ϕ_μ upon μ is not essential in proving Theorems 2.2, 2.3 as the only property of interest is that the derivatives of ϕ_μ are bounded from above and below uniformly over μ (which follows from our assumptions about Φ). This is the reason why in the next sections we will drop the explicit dependence on μ and work with a generic warping ϕ.

2.5.2 Proof of Theorem 2.2

We will assume without loss of generality that our curvelet φ_μ is centered near zero ($k = 0$) and is nearly vertical ($\theta_\mu = 0$).

Set $m_\mu = T_1\varphi_\mu$. We first show that m_μ obeys the smoothness and spatial localization estimate of a molecule (2.12). With the same notations as before, recall that m_μ is given by

$$m_\mu(x) = \int e^{ix\cdot\xi} b_\mu(x,\xi)\hat{\varphi}_\mu(\xi)\, d\xi, \quad b_\mu(x,\xi) = e^{i\delta(\phi^{-1}(x),\xi)}\sigma(\phi^{-1}(x),\xi). \tag{2.51}$$

To study the spatial decay of $m_\mu(x)$, we introduce the differential operator

$$L_\xi = I - 2^{2j}\frac{\partial^2}{\partial\xi_1^2} - 2^j\frac{\partial^2}{\partial\xi_2^2},$$

and evaluate the integral (2.51) using an integration by parts argument. First, observe that

$$L_\xi^N e^{ix\cdot\xi} = \left(1 + |2^j x_1|^2 + |2^{j/2}x_2|^2\right)^N e^{ix\cdot\xi}.$$

Second, we claim that for every integer $N \geq 0$,

$$|L_\xi^N[b_\mu(x,\xi)\hat{\varphi}_\mu(\xi)]| \leq C\cdot 2^{-3j/4}. \tag{2.52}$$

(The factor $2^{-3j/4}$ comes from the L^2 normalization of $\hat{\varphi}_\mu$.) This inequality is proved in appendix A.2. Hence,

$$m_\mu(x) = \left(1 + |2^j x_1|^2 + |2^{j/2}x_2|^2\right)^{-N} \int L_\xi^N[b_\mu(x,\xi)\hat{\varphi}_\mu(\xi)]\, e^{ix\cdot\xi}.$$

Since $|L_\xi^N[b_\mu(x,\xi)\hat{\varphi}_\mu(\xi)]| \leq C\cdot 2^{-3j/4}$ and is supported on a dyadic rectangle R_μ, of length about 2^j and width $2^{j/2}$, we then established that

$$|m_\mu(x)| \leq C\cdot\frac{2^{3j/4}}{\left(1 + |2^j x_1|^2 + |2^{j/2}x_2|^2\right)^N}.$$

The derivatives of m_μ are essentially treated in the same way. Begin with

$$\begin{aligned}\partial_x^\alpha(e^{ix\cdot\xi}b_\mu(x,\xi)) &= \sum_{\beta+\varphi\le\alpha}\partial^\beta(e^{ix\cdot\xi})\,\partial^\varphi(b_\mu(x,\xi))\\ &= \sum_{\beta+\varphi\le\alpha}\partial^\varphi(b_\mu(x,\xi))\,\xi^\beta e^{ix\cdot\xi}\end{aligned}$$

Therefore, the partial derivatives of m_μ are given by

$$(\partial_x^\alpha m_\mu)(x) = \sum_{\beta+\varphi\le\alpha} I_{\beta,\varphi}(x), \tag{2.53}$$

where

$$I_{\beta,\varphi}(x) = \int e^{ix\cdot\xi}\partial_x^\varphi(b_\mu(x,\xi))\xi^\beta\hat{\varphi}_\mu(\xi)\,d\xi. \tag{2.54}$$

First, observe that on the support of $\hat{\varphi}_\mu$, $|\xi|^\beta$ obeys $|\xi|^\beta \le C\cdot 2^{j\beta_1}\cdot 2^{j\beta_2/2}$. Second, the term $\partial_x^\varphi b(x,\xi)$ is of the same nature as $b_\mu(x,\xi)$ in the sense that it obeys all the same estimates as before. In particular, we claim that for every integer $N\ge 0$,

$$|L_\xi^N[\partial_x^\varphi b_\mu(x,\xi)\xi^\beta\hat{\varphi}_\mu(\xi)]| \le C\cdot 2^{-3j/4}\cdot 2^{j\beta_1}\cdot 2^{j\beta_2/2}. \tag{2.55}$$

Hence, the same argument as before gives

$$|I_{\beta,\varphi}(x)| \le C\cdot\frac{2^{3j/4}\cdot 2^{j\beta_1}\cdot 2^{j\beta_2/2}}{\left(1+|2^j x_1|^2+|2^{j/2}x_2|^2\right)^N}.$$

Now since $\beta\le\alpha$, we may conclude that

$$|(\partial_x^\alpha m_\mu)(x)| \le C\cdot\frac{2^{3j/4}\cdot 2^{j\alpha_1}\cdot 2^{j\alpha_2/2}}{\left(1+|2^j x_1|^2+|2^{j/2}x_2|^2\right)^N}.$$

This establishes the smoothness and localization property.

The above analysis shows that m_μ is a "ridge" of effective length $2^{-j/2}$ and width 2^{-j}; to prove that m_μ is a molecule, we now need to evidence its oscillatory behavior across the ridge. In other words, we are interested in the size of the Fourier transform at low frequencies (2.13)–(2.16).

Formally, the Fourier transform of m_μ is given by

$$\hat{m}_\mu(\xi) = \iint e^{ix\cdot(\eta-\xi)} b_\mu(x,\eta)\hat{\varphi}_\mu(\eta)\, dxd\eta. \tag{2.56}$$

We should point out that because the amplitude b is not of compact support in x, the sense in which (2.56) holds is not obvious. This is a well known phenomenon in Fourier analysis and a classical technique to circumvent such difficulties would be to multiply m_μ (or equivalently b_μ) by a smooth and compactly supported cut-off function $\chi(\epsilon x)$ and let ϵ tend to zero. We omit those details as they are standard.

Set $D_1 = -i\frac{\partial}{\partial x_1}$. To develop bounds on $|\hat{m}_\mu(\xi)|$, observe that

$$D_1^N e^{ix\cdot\eta} = (\eta_1)^N.$$

An integration by parts then gives

$$\hat{m}_\mu(\xi) = \iint e^{ix\cdot\eta} D_1^N \left(e^{-ix\cdot\xi}\, b_\mu(x,\xi) \right) \eta_1^{-N} \hat{\varphi}_\mu(\eta)\, dxd\eta.$$

Hence,

$$\hat{m}_\mu(\xi) = \sum_{m=0}^{N} c_m\, \xi_1^m \hat{F}_m(\xi),$$

where

$$F_m(x) = \int e^{ix\cdot\eta} (\partial_{x_1}^{n-m} b(x,\xi))\, \eta_1^{-n} \hat{\varphi}_\mu(\eta)\, d\eta.$$

Note that F_m is exactly of the same form as (2.54)—but with η_1^{-n} instead of η^β—and therefore, the exact same argument as before gives

$$|F_m(x)| \leq C \cdot \frac{2^{3j/4} \cdot 2^{-jn}}{\left(1 + |2^j x_1|^2 + |2^{j/2} x_2|^2\right)^N}.$$

We then established

$$\|\hat{F}_m\|_{L^\infty} \leq \|F\|_{L_1} \leq C_m \cdot 2^{-3j/4} \cdot 2^{-jn},$$

which gives

$$|m_\mu(\xi)| \leq C \cdot 2^{-3j/4} \cdot 2^{-jn} \cdot (1 + |\xi|^n),$$

as required. This finishes the proof of Theorem 2.2.

The careful reader will object that we did not study the case of coarse scale curvelets; it is obvious that coarse scale elements are mapped into coarse scale molecules and, here, the argument would not require the deployment of the sophisticated tools we exposed above. We omit the proof.

2.5.3 Atomic Decompositions

As we will see later, to prove our main result and especially Theorem 2.3, it would be most helpful to work with tight frames of curvelet compactly supported in space. Unfortunately, it is unclear at this point how to construct such tight frames with nice frequency localization properties. However, there exist useful atomic decompositions with compactly supported curvelet-like atoms. We now explore such decompositions.

In this section, the notation $f_{a,\theta}$ refers to the function obtained from f after applying a parabolic scaling and a rotation

$$f_{a,\theta}(x) = a^{-3/4} f\left(D_a R_\theta x\right), \quad D_a = \begin{pmatrix} 1/a & 0 \\ 0 & 1/\sqrt{a} \end{pmatrix},$$

and where R_θ is the rotation matrix which maps the vector $(1,0)$ into $(\cos\theta, -\sin\theta)$. Note that this is an isometry as

$$\|f_{a,\theta}\|_{L_2} = \|f\|_{L_2}.$$

In [73], Smith proved the following result: let $\tilde{\psi}$ be a Schwartz function obeying $\hat{\tilde{\psi}}(1,0) \neq 0$; then one can find another Schwartz function ψ, and a function $q(\xi)$ such that the following formula holds

$$q(\xi) \int_{a \leq 1} \hat{\tilde{\psi}}_{a,\theta}(\xi) \hat{\psi}_{a,\theta}(\xi)\, a da d\theta = r(\xi); \tag{2.57}$$

here r is a smooth cut-off function obeying

$$r(\xi) = \begin{cases} 1 & |\xi| \geq 2 \\ 0 & |\xi| \leq 1 \end{cases},$$

and q is a standard Fourier multiplier of order zero; that is, for each multiindex α, there exists a constant C_α such that

$$|\partial_\xi^\alpha q(\xi)| \leq C_\alpha (1 + |\xi|)^{-|\alpha|}.$$

This formula is useful because it allows us to express any object whose Fourier transform vanishes on $\{|\xi| \leq 2\}$ as a continuous superposition of curvelet-like elements. We now make some specific choices for φ. In the remainder of this section, we will take $\tilde{\psi}(x) = \psi(-x)$ and the function ψ of the form

$$\psi(x_1, x_2) = \psi^D(x_1)\, \varphi(x_2), \tag{2.58}$$

where both φ and ψ^D are compactly supported and obey

$$\operatorname{Supp} \varphi \subset [0,1], \quad \operatorname{Supp} \psi^D \subset [0,1].$$

We will assume that φ and ψ^D are C^∞ and that the function ψ^D has vanishing moments up to order D, i.e.,

$$\int \psi^D(x_1) x_1^k \, dx_1 = 0, \quad k = 0, 1, \ldots, D. \tag{2.59}$$

For each $a \leq 1$, each $b \in \mathbb{R}^2$ and each $\theta \in [0, 2\pi)$, introduce

$$\psi_{a,\theta,b}(x) := \psi_{a,\theta}(x - b) = a^{-3/4} \psi\left(D_a R_\theta (x - b)\right); \tag{2.60}$$

and given an object f, define coefficients by

$$CCT(f)(a, b, \theta) = \int \overline{\psi_{a,\theta,b}(x)} f(x) dx. \tag{2.61}$$

Now, suppose for instance that $\hat{f}$ vanishes over $|\xi| \leq 2$, then (2.57) gives the exact reconstruction formula

$$f(x) = \int_{a \geq 1} CCT(q(D)f)(a, b, \theta) \psi_{a,\theta,b}(x) \mu(da d\theta db), \tag{2.62}$$

with $\mu(da d\theta db) = a da d\theta db$. In the remainder of this section, we will use the shorter notation $d\mu$ for $\mu(da d\theta db)$.

As is now well established, the reproducing formula may be turned into a so-called "atomic decomposition." Not surprisingly, our atomic decomposition will just mimic the discretization of the curvelet frame as introduced in Section 2.2. With the notations of that section, we introduce the cells Q_μ defined as follows: for $j \geq 0$, $\ell = 0, 1, \ldots, 2^{\lfloor j/2 \rfloor} - 1$ and

$k = (k_1, k_2) \in \mathbb{Z}^2$, the cell Q_μ is the collections of triples (a, θ, b) for which

$$2^{-(j+1)} \le a < 2^{-j}, \quad |\theta - \theta_\mu| \le \frac{\pi}{2} 2^{-\lfloor j/2 \rfloor}$$

and

$$D_{2^{-j}} R_{\theta_\mu} b \in [k_1, k_1 + 1) \times [k_2, k_2 + 1).$$

Note that $\int_{Q_\mu} d\mu = 3\pi/2$ for j even, and 3π for j odd. We may then break the integral (2.62) into a sum of terms arising from different cells, namely,

$$f(x) = \sum_\mu \alpha_\mu \rho_\mu(x) \tag{2.63}$$

where

$$\begin{aligned} \alpha_\mu &= \|CCT(q(D)f)\|_{L_2(Q_\mu)}, \\ \rho_\mu(x) &= \frac{1}{\alpha_\mu} \int_{Q_\mu} CCT(q(D)f))(a, b, \theta) \psi_{a,\theta,b}(x)\, d\mu. \end{aligned} \tag{2.64}$$

Of course, the decomposition (2.63) greatly resembles the tight frame expansion, compare (1.9). In particular, the atoms ρ_μ are curvelet-like in the sense that they share all the properties of the tight frame $(\varphi_\mu)_\mu$ – only they are compactly supported in space. In the remainder of the chapter, we will call these elements *curvelet atoms.* Below are some crucial properties of these atoms. Please note that we are not talking about wave atoms here, but merely a specific type of curvelet-like waveforms which happen to be named "atoms" as well.

Lemma 2.7. *Rewrite the atoms ρ_μ as $\rho_\mu(x) = 2^{3j/4} a^{(\mu)} \left(D_{2^{-j}} R_{\theta_\mu} x - k\right)$. In other words, ρ_μ is obtained from $a^{(\mu)}$ after parabolic scaling, rotation, and translation. For all μ, the $a^{(\mu)}$ verify the following properties.*

- *Compact support;*

$$Supp\ a^{(\mu)} \subset cQ. \tag{2.65}$$

- *Nearly vanishing moment along the horizontal axis; let $m = D/2$. Then for each*

$k = 0, 1, \ldots, m$, *there is a constant* C_m *such that*

$$\int a^{(\mu)}(x_1, x_2) x_1^k \, dx_1 \leq C_m \cdot 2^{-j(m+1)}. \tag{2.66}$$

- *Regularity; for every multiindex* α

$$|\partial_x^\alpha a^{(\mu)}(x)| \leq C_\alpha. \tag{2.67}$$

In (2.66) and (2.67), the constants may be chosen independently of μ *and* f.

Proof. See appendix A.2 □

Needless to say that curvelet atoms are molecules with spatial compact support, compare lemma 2.7 with the definition of a molecule. Finally, observe (and this is important) that it is of course possible to decompose a molecule into a series of atoms

$$m_\mu = \sum_{\mu'} \alpha_{\mu\mu'} \rho_{\mu'}.$$

The coefficients would then obey the same estimate as in lemma 2.1

$$|\alpha_{\mu\mu'}| \leq C_N \cdot \omega(\mu, \mu')^{-N}, \tag{2.68}$$

and in particular, for each $p > 0$,

$$\sup_\mu \sum_{\mu'} |\alpha_{\mu\mu'}|^p < C_p.$$

This is briefly justified in appendix A.2.

2.5.4 Proof of Theorem 2.3

As mentioned earlier, curvelet atoms depend in a nonessential way upon the object f we wish to analyze and we shall drop this dependence in our notations. To prove Theorem 2.3, recall that we need to show that for each curvelet atom ρ_μ with regularity R, the "warped" atom $\rho_\mu \circ \phi$ is also a curvelet atom, with the same regularity.

As in Section 2.5.3, we suppose our curvelet atom is of the form

$$\rho_\mu(x) = 2^{3j/4} a^{(\mu)}(D_{2^{-j}} R_{\theta_\mu}(x - x_\mu)),$$

where $a^{(\mu)}$ obeys the conditions of Lemma 2.7. (Here, the location x_μ may be formally defined by $x_\mu = (D_{2^{-j}} R_{\theta_\mu})^{-1} k_\delta$.) Define y_μ and A_μ by

$$y_\mu = \phi^{-1}(x_\mu), \quad \text{and} \quad A_\mu = (\nabla\phi)(y_\mu) \tag{2.69}$$

so that

$$\phi(y) = x_\mu + A_\mu(y - y_\mu) + g(y - y_\mu).$$

With these notations, it is clear that the warped atom $\rho_\mu \circ \phi$ will be centered near the point y_μ; that is,

$$\rho_\mu(\phi(y)) = 2^{3j/4} a^{(\mu)} \left(D_{2^{-j}} R_{\theta_\mu}(A_\mu(y - y_\mu) + g(y - y_\mu)) \right).$$

To simplify matters, we first assume that A_μ is the identity and show that $\rho_\mu \circ \phi$ is a curvelet atom with the same scale and orientation as ρ_μ. Later, we will see that in general, $\rho_\mu \circ \phi$ is an atom whose orientation depends upon A_μ, and whose scale may be taken to be the same as that of ρ_μ. Assume without loss of generality that $\theta_\mu = 0$ and $y_\mu = 0$ (statements for arbitrary orientations and locations are obtained in an obvious fashion) so that

$$\rho_\mu(\phi(y)) = 2^{3j/4} a^{(\mu)} \left(D_{2^{-j}}(y + g(y)) \right) = 2^{3j/4} b^{(\mu)}(D_{2^{-j}} y), \tag{2.70}$$

with

$$b^{(\mu)}(y) = a^{(\mu)} \left(y + D_{2^{-j}} g(D_{2^j} y) \right).$$

The atom $a^{(\mu)}$ is supported over a square of sidelength about 1; likewise, $b^{(\mu)}$ is also compactly supported in a box of roughly the same size—uniformly over μ. We then need to derive smoothness estimates and show that $b^{(\mu)}$ obeys

$$|\partial^\alpha b^{(\mu)}(y)| \le C_\alpha, \quad |\alpha| \le R. \tag{2.71}$$

Over the support of $\rho_\mu \circ \phi$, $g = (g_1, g_2)$ deviates little from zero and for each $k = 1, 2$, g_k

obeys

$$|g_k(y)| \le C \cdot 2^{-j}, \quad |\partial^\alpha g_k(y)| \le C \cdot 2^{-j/2}, \ |\alpha| = 1.$$

Similarly, for each α, $|\alpha| > 1$,

$$|\partial^\alpha g_k(y)| \le C_\alpha. \tag{2.72}$$

These estimates hold uniformly over μ. It follows that for $|y_1|, |y_2| \le C$ and each α, the perturbation g obeys

$$2^j \cdot |\partial^\alpha g_1(2^{-j}y_1, 2^{-j/2}y_2)| \le C_\alpha, \quad 2^{j/2} \cdot |\partial^\alpha g_2(2^{-j}y_1, 2^{-j/2}y_2)| \le C_\alpha. \tag{2.73}$$

The bound (2.71) is then a simple consequence of (2.73) together with the fact that all the derivatives of $a^{(\mu)}$ up to order R are bounded, uniformly over μ.

We now show that $\rho_\mu \circ \phi$ exhibits the appropriate behavior at low frequencies.

$$\begin{aligned}\widehat{\rho_\mu \circ \phi}(\xi) &= \int e^{-ix\cdot\xi} \rho_\mu(\phi(x))\, dx \\ &= \int e^{-i\phi^{-1}(x)\cdot\xi} \rho_\mu(x) \frac{dx}{|\det \nabla\phi|(\phi^{-1}(x))}.\end{aligned}$$

We will use the nearly vanishing moment property of ρ_μ. Set

$$S_\xi(x) = e^{-i\phi^{-1}(x)\cdot\xi}/|\det \nabla\phi|(\phi^{-1}(x));$$

note that over the support of ρ_μ and for each $N \le R$, we have available the following upper bound on the partial derivative of S_ξ

$$|\partial_1^N S_\xi(x)| \le C_N \cdot (1 + |\xi|)^N.$$

Classical arguments give

$$\widehat{\rho_\mu \circ \phi}(\xi) = \sum_{k=0}^{n-1} \int \frac{\partial_1^k S_\xi(0, x_2)}{k!}\, dx_2 \int \rho_\mu(x_1, x_2) x_1^k\, dx_1 dx_2 + E, \tag{2.74}$$

where E is a remainder term obeying

$$|E| \le C_n \cdot 2^{-3j/4} \cdot 2^{-jn} \cdot \sup |\partial_1^n S_\xi(x)| \le C_n \cdot 2^{-3j/4} \cdot 2^{-jn}(1 + |\xi|^n). \tag{2.75}$$

The near-vanishing moment property gives that each term in the right-hand side of (2.74) obeys the estimate in (2.75). This proves that the Fourier transform of $\rho_\mu \circ \phi$ obeys

$$|\widehat{\rho_\mu \circ \phi}(\xi)| \le C_n \cdot 2^{-jn}(1 + |\xi|^n)$$

as required.

We now discuss the case where the matrix A_μ is not the identity. In this case, (2.70) becomes

$$\rho_\mu(\phi(y)) = m_\mu(A_\mu y),$$

with

$$m_\mu(y) = 2^{3j/4} a^{(\mu)} \left(D_{2^j}(y + \tilde{g}(y))\right), \quad \text{and} \quad \tilde{g}(y) = g(A_\mu^{-1} y).$$

Our assumptions about FIOs guarantee that $|A_\mu^{-1}|$ is uniformly bounded and, therefore, it follows from the previous analysis that m_μ is a curvelet atom. As a consequence $\rho_\mu \circ \phi$ is a curvelet atom with the same regularity R since it is clear that bounded linear transformations of the plane map curvelet atoms into curvelet atoms.

2.5.5 Proof of Theorem 2.1

Let φ_{μ_0} be a fixed curvelet and decompose T as $T_{2,\mu_0} \circ T_{1,\mu_0}$. First, Theorem 2.2 proved that $T_{1,\mu_0}\varphi_{\mu_0}$ is a curvelet molecule m_{μ_0} which we will express as a superposition of curvelet atoms ρ_{μ_1}

$$T_{1,\mu_0}\varphi_{\mu_0} = m_{\mu_0} = \sum_{\mu_1} \beta_0(\mu_1, \mu_0)\rho_{\mu_1}.$$

Second, for each μ_1, Theorem 2.3 shows that $T_{2,\mu_1}\rho_{\mu_1}$ is a molecule $m_{h(\mu_1)}$ at the location $h(\mu_1)$. We are not exactly in that setting since in $T_{2,\mu_0}\rho_{\mu_1}$, the subscripts do not, in general, match. This does not pose any difficulty since Theorem 2.3 can be understood as a statement concerning general warpings ϕ. We can define the map h_{μ_0} as induced by the transformation $(x, \xi) \to (y, \eta)$ given by

$$x = \nabla_x \Phi(y, \xi_{\mu_0}), \qquad \eta = \nabla_\xi(y, \xi_{\mu_0})$$

(compare this with equation (2.47)). Then, according to Theorem 2.3, $T_{2,\mu_0}\rho_{\mu_1}$ is a molecule $m_{h_{\mu_0}(\mu_1)}$ at the location $h_{\mu_0}(\mu_1)$. So

$$\langle \varphi_{\mu_2}, T_{2,\mu_0}\rho_{\mu_1}\rangle = \beta_1(\mu_2, h_{\mu_0}(\mu_1)).$$

Hence,

$$\langle \varphi_{\mu_2}, T\varphi_{\mu_0}\rangle = \sum_{\mu_1} \beta_1(\mu_2, h_{\mu_0}(\mu_1))\beta_0(\mu_1, \mu_0).$$

Of course, both β_0 and β_1 obey very special decay properties.

- By Theorem 2.2 and Lemma 2.1, $|\beta_0(\mu_1, \mu_0)| \leq C_n \cdot \omega(\mu_1, \mu_0)^{-N}$ for arbitrarily large $N > 0$, provided that the selected atoms are regular enough.

- By Theorem 2.3 and Lemma 2.1, $|\beta_1(\mu_2, h_{\mu_0}(\mu_1))| \leq C_N \cdot \omega(\mu_2, h_{\mu_0}(\mu_1))^{-N}$ for arbitrarily large $N > 0$, provided that the selected atoms are regular enough.

Theorem 2.1 now follows from the observation that

$$\sum_{\mu_1} \omega(\mu_2, h_{\mu_0}(\mu_1))^{-N} \cdot \omega(\mu_1, \mu_0)^{-N} \leq C_n \cdot \omega(\mu_2, h_{\mu_0}(\mu_0))^{-(N-1)}, \tag{2.76}$$

This is an immediate consequence of properties 3 and 4 of the pseudodistance ω, see proposition 2.1.

Cases involving coarse scale elements are treated similarly and we omit the proof. The boundedness from ℓ^p to ℓ^p for every $p > 0$ follows from the same argument as in the proof of Theorem 1.1.

2.6 Discussion

All along we specialized our discussion to the special case where the dimension of the spatial variable is $n = 2$. It is clear that nothing in our arguments depends upon this specific assumption. Indeed, we could just as well construct tight frames of curvelets in arbitrary dimensions by smoothly partitioning the frequency plane into dyadic coronae, which would then be angularly localized near regions of sidelength length 2^j in the radial direction and $2^{j/2}$ in all the other directions; in order to this, we would use smooth partitions of the unit sphere of $\mathbb{R}^n$ into spherical caps of radius about $2^{-j/2}$. All of our analysis would apply as

is, and would prove versions of Theorem 1.1 in arbitrary dimensions.

Our main result assumes that the coefficients of the equation (1.7) be smooth. In many applications of interest, however, the coefficients may be smooth away from singular smooth surfaces. In geophysics for example, we typically have different layers with very different physical properties. A very important question would be to know how our analysis would adapt to this situation. In fact, it seems natural to believe that sparsity would continue to hold in this more general setting. Intuitively, the wave group would still be approximated by rigid motion along the Hamiltonian flow. Only, one would need to account for possible reflections/refractions. A curvelet hitting a singularity at a small angle of incidence would typically produce two curvelets, a reflected and a refracted curvelet. This is merely an intuition which one would need to justify by a careful analysis quantifying the behavior of a curvelet near the interface (here, the singular surface). We regard this type of question as an important extension to this work.

Chapter 3

Fast Discrete Curvelet Transforms

This chapter is organized as follows. Section 3.1 introduces the main ideas underlying the wrapping-based digital implementation of curvelets. Its mathematical properties are then detailed in Section 3.2. Section 3.3 discusses refinements and extensions of the ideas underlying the discrete transforms while Section 3.4 illustrates our methods with a few numerical experiments. Finally, we conclude with Section 3.5 which explains connections with the work of others, and outlines possible applications of these transforms. A few open problems are listed in the conclusion (Chapter 5).

The software package CurveLab implements the transforms proposed in this paper, and is available at http://www.curvelet.org. It contains the Matlab and C++ implementations of both the USFFT-based [16] and the wrapping-based transforms. Several Matlab scripts are provided to demonstrate how to use this software. Additionally, three different implementations of the 3D discrete curvelet transform are also included.

3.1 Digital Curvelet Transforms

In this chapter, we propose an implementation of the curvelet transform which is faithful to the mathematical transformation outlined in section 2.2. These digital transformations are linear and take as input Cartesian arrays of the form $f[t_1, t_2]$, $0 \le t_1, t_2 < n$, which allows us to think of the output as a collection of coefficients $c^D(j, \ell, k)$ obtained by the digital analog to (2.8)

$$c^D(j, \ell, k) := \sum_{0 \le t_1, t_2 < n} f[t_1, t_2] \, \overline{\varphi^D_{j,\ell,k}[t_1, t_2]}, \tag{3.1}$$

where each $\varphi^D_{j,\ell,k}$ is a digital curvelet waveform (here and below, the superscript D stands for "digital"). As is standard in scientific computations, we will actually never build these digital waveforms which are implicitly defined by the algorithms; formally, they are the rows of the matrix representing the linear transformation and are also known as Riesz representers. We merely introduce these waveforms because it will make the exposition clearer and because it provides a useful way to explain the relationship with the continuous-time transformation.

Let us now introduce the architecture of the curvelet transform.

3.1.1 Digital Coronization

In the continuous-time definition (2.7), the window U_j smoothly extracts frequencies near the dyadic corona $\{2^j \le r \le 2^{j+1}\}$ and near the angle $\{-\pi \cdot 2^{-j/2} \le \theta \le \pi \cdot 2^{-j/2}\}$. Coronae and rotations are not especially adapted to Cartesian arrays. Instead, it is convenient to replace these concepts by Cartesian equivalents; here, "Cartesian coronae" based on concentric squares (instead of circles) and shears. For example, the Cartesian analog to the family $(W_j)_{j\ge 0}$, $W_j(\omega) = W(2^{-j}\omega)$, would be a window of the form

$$\tilde{W}_j(\omega) = \sqrt{\Phi^2_{j+1}(\omega) - \Phi^2_j(\omega)}, \quad j \ge 0,$$

where Φ is defined as the product of low-pass one dimensional windows

$$\Phi_j(\omega_1, \omega_2) = \phi(2^{-j}\omega_1)\, \phi(2^{-j}\omega_2).$$

The function ϕ obeys $0 \le \phi \le 1$, might be equal to 1 on $[-1/2, 1/2]$, and vanishes outside of $[-2, 2]$. It is immediate to check that

$$\Phi_0(\omega)^2 + \sum_{j\ge 0} \tilde{W}^2_j(\omega) = 1. \tag{3.2}$$

We have just seen how to separate scales in a Cartesian-friendly fashion and now examine the angular localization. Suppose that V is as before, i.e., obeys (2.6) and set

$$V_j(\omega) = V(2^{\lfloor j/2 \rfloor}\omega_2/\omega_1).$$

We can then use $\tilde{W}_j$ and V_j to define the "Cartesian" window

$$\tilde{U}_j(\omega) := \tilde{W}_j(\omega) V_j(\omega). \tag{3.3}$$

It is clear that $\tilde{U}_j$ isolates frequencies near the wedge $\{(\omega_1, \omega_2) : 2^j \le \omega_1 \le 2^{j+1},\ -2^{-j/2} \le \omega_2/\omega_1 \le 2^{-j/2}\}$, and is a Cartesian equivalent to the "polar" window of Section 2.2. Introduce now the set of equispaced slopes $\tan\theta_\ell := \ell \cdot 2^{-\lfloor j/2 \rfloor}$, $\ell = -2^{\lfloor j/2 \rfloor}, \ldots, 2^{\lfloor j/2 \rfloor} - 1$, and define

$$\tilde{U}_{j,\ell}(\omega) := W_j(\omega) V_j(S_{\theta_\ell}\, \omega),$$

where S_θ is the shear matrix,

$$S_\theta := \begin{pmatrix} 1 & 0 \\ -\tan\theta & 1 \end{pmatrix}.$$

The angles θ_ℓ are not equispaced here but the slopes are. When completed by symmetry around the origin and rotation by $\pm\pi/2$ radians, the $\tilde{U}_{j,\ell}$ define the Cartesian analog to the family $U_j(R_{\theta_\ell}\omega)$ of Section 2.2. The family $\tilde{U}_{j,\ell}$ implies a concentric tiling whose geometry is pictured in Figure 3.1.[1]

By construction, $V_j(S_{\theta_\ell}\, \omega) = V(2^{\lfloor j/2 \rfloor}\omega_2/\omega_1 - \ell)$ and for each $\omega = (\omega_1, \omega_2)$ with $\omega_1 > 0$, say, (2.6) gives

$$\sum_{\ell=-\infty}^{\infty} |V_j(S_{\theta_\ell}\, \omega)|^2 = 1.$$

Because of the support constraint on the function V, the above sum restricted to the angles of interest, $-1 \le \tan\theta_\ell < 1$, obeys $\sum_{\text{all angles}} |V_j(S_{\theta_\ell}\, \omega)|^2 = 1$, for $\omega_2/\omega_1 \in [-1 + 2^{-\lfloor j/2 \rfloor}, 1 - 2^{-\lfloor j/2 \rfloor}]$. Therefore, it follows from (3.2) that

$$\sum_{\text{all scales}} \sum_{\text{all angles}} |\tilde{U}_{j,\ell}(\omega)|^2 = 1. \tag{3.5}$$

[1] There are other ways of defining such localizing windows. An alternative might be to select $\tilde{U}_j$ as

$$\tilde{U}_j(\omega) := \psi_j(\omega_1) V_j(\omega), \tag{3.4}$$

where $\psi_j(\omega_1) = \psi(2^{-j}\omega_1)$ with $\psi(\omega_1) = \sqrt{\phi(\omega_1/2)^2 - \phi(\omega_1)^2}$ a bandpass profile, and to define for each $\theta_\ell \in [-\pi/4, \pi/4)$

$$\tilde{U}_{j,\ell}(\omega) := \psi_j(\omega_1) V_j(S_{\theta_\ell}\, \omega) = \tilde{U}_j(S_{\theta_\ell}\, \omega).$$

With this special definition, the windows are shear-invariant at any given scale. In practice, both these choices are almost equivalent since for a large number of angles of interest, many ϕ would actually give identical windows $\tilde{U}_{j,\ell}$.

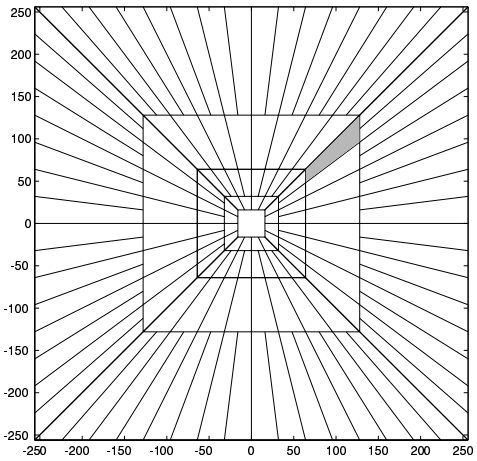

Figure 3.1: The figure illustrates the basic digital tiling. The windows $\tilde{U}_{j,\ell}$ smoothly localize the Fourier transform near the sheared wedges obeying the parabolic scaling. The shaded region represents one such typical wedge.

There is a way to define "corner" windows specially adapted to junctions over the four quadrants (east, south, west, north) so that (3.5) holds for every $\omega \in \mathbb{R}^2$. We postpone this technical issue to Section 3.3.2.

The pseudo-polar tiling of the frequency plane with trapezoids, in Figure 3.1, is already well-established as a data-friendly alternative to the ideal polar tiling. It was perhaps first introduced in two articles that appeared as book chapters in the same book, *Beyond Wavelets*, in 2003. The first construction is that of *contourlets* [31] and is based on a cascade of properly sheared directional filters. On the other hand, *ridgelet packets* [40] are defined directly in the frequency plane via interpolation onto a pseudo-polar grid aligned with the trapezoids.

In the next section we explain the rest of the architecture, based on the *wrapping* operation. Candès and Donoho proposed another implementation, via USFFT, which is described in [16]. In a nutshell, the two implementations differ in the way curvelets at a given scale and angle are translated with respect to each other. In the USFFT-based version the translation grid is tilted to be aligned with the orientation of the curvelet, yielding the most faithful discretization of the continuous definition. In the Wrapping version the grid is the same for every angle within each quadrant – yet each curvelet is given the proper orientation. As a result, the wrapping-based transform may be simpler to understand and implement.

3.1.2 Digital Curvelet Transform via Wrapping

The "wrapping" approach assumes a regular rectangular grid to translate curvelets at each scale and angle, and defines "Cartesian" curvelets as

$$c(j,\ell,k) = \int \hat{f}(\omega)\tilde{U}_j(S_{\theta_\ell}^{-1}\,\omega)e^{i\langle b,\omega\rangle}\,d\omega. \tag{3.6}$$

In the above formula we will take $b \simeq (k_1 2^{-j}, k_2 2^{-j/2})$, taking on values on a rectangular grid. This formula for b is understood when $\theta \in (-\frac{\pi}{4}, \frac{\pi}{4})$ or $(\frac{3\pi}{4}, \frac{5\pi}{4})$, otherwise the roles of $L_{1,j}$ and $L_{2,j}$ are to be exchanged.

The difficulty behind this approach is that, in the frequency plane, the window $\tilde{U}_{j,\ell}[n_1, n_2]$ does not fit in a rectangle of size $\sim 2^j \times 2^{j/2}$, aligned with the axes, in which the 2D IFFT could be applied to compute (3.6). After discretization, the integral over ω becomes a sum over n_1, n_2 which would extend beyond the bounds allowed by the 2D IFFT. The resemblance of (3.6) with a standard 2D inverse FFT is, in that respect, only formal.

To understand why respecting rectangle sizes is a concern, we recall that $\tilde{U}_{j,\ell}$ is supported in the parallelepipedal region

$$\mathcal{P}_{j,\ell} = S_{\theta_\ell}\,\mathcal{P}_j.$$

For most values of the angular variable θ_ℓ, $\mathcal{P}_{j,\ell}$ is supported inside a rectangle $\mathcal{R}_{j,\ell}$ aligned with the axes, and with sidelengths both on the order of 2^j. One could in principle use the 2D inverse FFT on this larger rectangle instead. This is close in spirit to the discretization of the continuous directional wavelet transform proposed by Vandergheynst and Gobbers in [86]. This seems ideal, but there is an apparent downside to this approach: dramatic oversampling of the coefficients. In other words, whereas the previous approach showed that it was possible to design curvelets with anisotropic spatial spacing of about $n/2^j$ in one direction and $n/2^{j/2}$ in the other, this approach would seem to require a naive regular rectangular grid with sidelength about $n/2^j$ in both directions. In other words, one would need to compute on the order of 2^{2j} coefficients per scale and angle as opposed, to only about $2^{3j/2}$ in the USFFT-based implementation. By looking at fine scale curvelets such that $2^j \asymp n$, this approach would require $O(n^{2.5})$ storage versus $O(n^2)$ for the USFFT version.

It is possible, however, to downsample the naive grid, and obtain for each scale and angle

a subgrid which has the same cardinality as that in use in the USFFT implementation. The idea is to periodize the frequency samples as we now explain.

As before, we let $\mathcal{P}_{j,\ell}$ be a parallelogram containing the support of the discrete localizing window $\tilde{U}_{j,\ell}[n_1, n_2]$. We suppose that at each scale j, there exist two constants $L_{1,j} \sim 2^j$ and $L_{2,j} \sim 2^{j/2}$ such that, for every orientation θ_ℓ, one can tile the two-dimensional plane with translates of $\mathcal{P}_{j,\ell}$ by multiples of $L_{1,j}$ in the horizontal direction and $L_{2,j}$ in the vertical direction. The corresponding periodization of the windowed data $d[n_1, n_2] = \tilde{U}_{j,\ell}[n_1, n_2]\hat{f}[n_1, n_2]$ reads

$$Wd[n_1, n_2] = \sum_{m_1 \in \mathbb{Z}} \sum_{m_2 \in Z} d[n_1 + m_1 L_{1,j}, n_2 + m_2 L_{2,j}]$$

The *wrapped* windowed data, around the origin, is then defined as the restriction of $Wd[n_1, n_2]$ to indices n_1, n_2 inside a rectangle with sides of length $L_{1,j} \times L_{2,j}$ near the origin:

$$0 \le n_1 < L_{1,j}, \qquad 0 \le n_2 < L_{2,j}.$$

Given indices (n_1, n_2) originally inside $\mathcal{P}_{j,\ell}$ (possibly much larger than $L_{1,j}$, $L_{2,j}$), the correspondence between the wrapped and the original indices is one-to-one. Hence, the wrapping transformation is a simple reindexing of the data. It is possible to express the wrapping of the array $d[n_1, n_2]$ around the origin even more simply by using the "modulo" function:

$$Wd[n_1 \bmod L_{1,j},\, n_2 \bmod L_{2,j}] = d[n_1, n_2], \tag{3.7}$$

with $(n_1, n_2) \in \mathcal{P}_{j,\ell}$. Intuitively, the modulo operation maps the original (n_1, n_2) into their new position near the origin.

For those angles in the range $\theta \in (\pi/4, 3\pi/4)$, the wrapping is similar, after exchanging the role of the coordinate axes. This is the situation shown in figure 3.2.

Equipped with this definition, the architecture of the FDCT via wrapping is as follows:

1. Apply the 2D FFT and obtain Fourier samples $\hat{f}[n_1, n_2]$, $-n/2 \le n_1, n_2 < n/2$.
2. For each scale j and angle ℓ, form the product $\tilde{U}_{j,\ell}[n_1, n_2]\hat{f}[n_1, n_2]$.

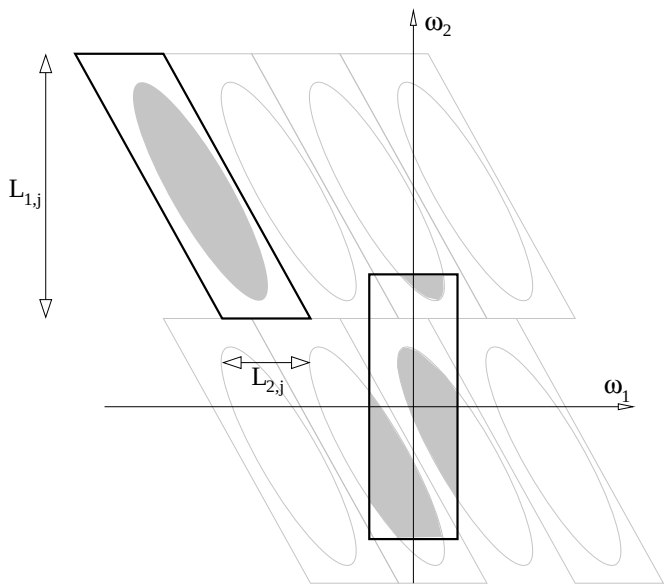

Figure 3.2: Wrapping data, intially inside a parallelogram, into a rectangle by periodicity. The angle θ is here in the range $(\pi/4, 3\pi/4)$. The black parallelogram is the tile $\mathcal{P}_{j,\ell}$ which contains the frequency support of the curvelet, whereas the gray parallelograms are the replicas resulting from periodization. The rectangle is centered at the origin. The wrapped ellipse appears "broken into pieces" but as we shall see, this is not an issue in the *periodic* rectangle, where the opposite edges are identified.

3. Wrap this product around the origin and obtain

$$\tilde{f}_{j,\ell}[n_1, n_2] = W(\tilde{U}_{j,\ell}\hat{f})[n_1, n_2],$$

   where the range for n_1 and n_2 is now $0 \le n_1 < L_{1,j}$ and $0 \le n_2 < L_{2,j}$ (for θ in the range $(-\pi/4, \pi/4)$).

4. Apply the inverse 2D FFT to each $\tilde{f}_{j,\ell}$, hence collecting the discrete coefficients $c^D(j, \ell, k)$.

It is clear that this algorithm has computational complexity $O(n^2 \log n)$ and in practice, its computational cost does not exceed that of 6 to 10 two-dimensional FFTs, see Section 3.4 for typical values of CPU times. In Section 3.2, we will detail some of the properties of this transform, namely, (1) it is an isometry, hence the inverse transform can simply be computed as the adjoint, and (2) it is faithful to the continuous transform.

3.1.3 FDCT Architecture

We finally close this section by listing the key elements of the implementation.

1. Frequency space is divided into dyadic annuli based on concentric squares.
2. Each annulus is subdivided into trapezoidal regions.
3. The wrapping operation uses extension by periodization to localize the Fourier samples in a rectangular region in which the IFFT can be applied. For a given scale, this corresponds only to two Cartesian sampling grids, one for all angles in the east-west quadrants, and one for the north-south quadrants.
4. Both forward transforms are specified in closed form, and are invertible, with inverse in closed form.
5. The design of appropriate digital curvelets at the finest scale, or outermost dyadic corona, is not straightforward because of boundary/periodicity issues. Possible solutions at the finest scale are discussed in Section 3.3.

3.2 FDCT via Frequency Wrapping

3.2.1 Riesz Representers

The naive technique suggested in Section 3.1 to obtain oversampled curvelet coefficients consists of a simple 2D inverse FFT, which reads

$$c^{D,O}(j,\ell,k) = \frac{1}{n^2} \sum_{n_1,n_2 \in \mathcal{R}_{j,\ell}} \hat{f}[n_1,n_2]\tilde{U}_{j,\ell}[n_1,n_2]e^{2\pi i(k_1 n_1/R_{1,j}+k_2 n_2/R_{2,j})}. \tag{3.8}$$

The superscripts D, O stand for Digital, Oversampled. As before, $\mathcal{R}_{j,\ell}$ is a rectangle of size $R_{1,j} \times R_{2,j}$, aligned with the Cartesian axes, and containing the parallelogram $\mathcal{P}_{j,\ell}$. Assume that $R_{1,j}$, $R_{2,j}$ divide the image size n. Then it is not hard to see that the coefficients $c^{D,O}(j,\ell,k)$ come from the discrete convolution of a curvelet with the signal $f(t_1,t_2)$, downsampled regularly in the sense that one selects only one out of every $n/R_{1,j} \times n/R_{2,j}$ pixel.

In general the dimensions $R_{1,j}$, $R_{2,j}$ of the rectangle are too large, as explained earlier. Equivalently, one wishes to downsample the convolution further. The idea of the wrapping

approach is to replace $R_{1,j}$ and $R_{2,j}$ in equation (3.8) by $L_{1,j}$ and $L_{2,j}$, the original dimensions of the parallelogram $\mathcal{P}_{j,\ell}$. In order to fit $\mathcal{P}_{j,\ell}$ into a rectangle with the same dimensions, we need to copy the data by periodicity, or wrap-around, as illustrated in Figure 3.2. This is just a relabeling of the frequency samples, of the form

$$n_1' = n_1 + m_1 L_{1,j}, \qquad n_2' = n_2 + m_2 L_{2,j},$$

for some adequate integers m_1 and m_2 themselves depending on n_1 and n_2.

The 2D inverse FFT of the wrapped array therefore reads

$$c^D(j,\ell,k) = \frac{1}{n^2} \sum_{n_1=0}^{L_{1,j}-1} \sum_{n_2=0}^{L_{2,j}-1} W(\tilde{U}_{j,\ell}\hat{f})[n_1,n_2] e^{2\pi i (k_1 n_1/L_{1,j} + k_2 n_2/L_{2,j})}. \tag{3.9}$$

Notice that the wrapping relabeling leaves the phase factors unchanged in the above formula, so we can also write it as[2]

$$c^D(j,\ell,k) = \frac{1}{n^2} \sum_{n_1=-n/2}^{n/2-1} \sum_{n_2=-n/2}^{n/2-1} \tilde{U}_{j,\ell}[n_1,n_2] \hat{f}[n_1,n_2] e^{2\pi i (k_1 n_1/L_{1,j} + k_2 n_2/L_{2,j})}.$$

It is then easy to conclude that we have correctly downsampled the convolution of f with the discrete curvelet, this time at every other $n/L_{1,j} \times n/L_{2,j}$ pixels. The following statement establishes precisely this fact, i.e., that the curvelet transform computed by wrapping is as geometrically faithful to the continuous transform as the sampling on the grid allows.

Proposition 3.1. *Let $\varphi^D_{j,\ell}$ be the "mother curvelet" at scale j and angle ℓ,*

$$\varphi^D_{j,\ell}(x) = \frac{1}{(2\pi)^2} \int e^{i\langle x,\omega\rangle} \tilde{U}_{j,\ell}(\omega)\, d\omega,$$

and $\varphi^\sharp_{j,\ell}$ denote its periodization over the unit square $[0,1]^2$,

$$\varphi^\sharp_{j,\ell}(x_1,x_2) = \sum_{m_1\in\mathbb{Z}} \sum_{m_2\in\mathbb{Z}} \varphi^D_{j,\ell}(x_1+m_1, x_2+m_2).$$

[2]The leading factor $\frac{1}{n^2}$ is not the standard one for the inverse FFT (that would be $\frac{1}{L_{1,j}L_{2,j}}$), but this choice of normalization is useful in the formulation of proposition 3.1. Yet another choice of normalization will be made later to make the transform an isometry.

In exact arithmetic, the coefficients in the East and West quadrants are given by

$$c^D(j,\ell,k) = \frac{1}{n^2}\sum_{t_1=0}^{n-1}\sum_{t_2=0}^{n-1} f[t_1,t_2]\,\overline{\varphi^{\sharp}_{j,\ell}}(\frac{t_1}{n}-\frac{k_1}{L_{1,j}},\frac{t_2}{n}-\frac{k_2}{L_{2,j}}). \tag{3.10}$$

This is a discrete circular convolution if and only if $L_{1,j}$ and $L_{2,j}$ both divide n. For angles in the North and South quadrants, reverse the roles of $L_{1,j}$ and $L_{2,j}$.

Proof. By definition, the East and West coefficients are given by the formula

$$c^D(j,\ell,k) = \frac{1}{n^2}\sum_{n_1=0}^{L_{1,j}-1}\sum_{n_2=0}^{L_{2,j}-1} e^{2\pi i k_1 n_1/L_{1,j}} e^{2\pi i k_2 n_2/L_{2,j}} W(\tilde{U}_{j,\ell}\hat{f})[n_1,n_2].$$

Let us change n_1 and n_2 to $n_1' = n_1 + m_1 L_{1,j}$, $n_2' = n_2 + m_2 L_{2,j}$, for appropriate integers m_1, m_2 (themselves depending on n_1 and n_2) so that $(2\pi n_1', 2\pi n_2') \in \mathcal{P}_{j,\ell}$, or more concisely, "$n_1', n_2'$ in tile." This is the unwrapping transformation, and leaves the phase factors unchanged. Notice that $n_1 = n_1' \bmod L_{1,j}$ and $n_2 = n_2' \bmod L_{2,j}$. We can then use the definition of wrapping in equation (3.7) to rewrite

$$c^D(j,\ell,k) = \frac{1}{n^2}\sum_{n_1,n_2 \text{ in tile}} e^{2\pi i k_1 n_1/L_{1,j}} e^{2\pi i k_2 n_2/L_{2,j}}\,\tilde{U}_{j,\ell}[n_1,n_2]\,\hat{f}[n_1,n_2].$$

We recall that the index-to-sample correspondence in the frequency plane is just

$$\tilde{U}_{j,\ell}[n_1,n_2] = \tilde{U}_{j,\ell}(2\pi n_1, 2\pi n_2).$$

It is also valid for $\hat{f}$, if we introduce $\hat{f}(\omega_1,\omega_2)$ as the trigonometric interpolant of the array $\hat{f}[n_1,n_2]$. Notice in passing that $\hat{f}(\omega_1,\omega_2)$ is periodic in ω outside of the fundamental cell, so we actually have

$$\hat{f}(2\pi n_1, 2\pi n_2) = \hat{f}[(n_1+\frac{n}{2}) \bmod n - \frac{n}{2}, (n_2+\frac{n}{2}) \bmod n - \frac{n}{2}] \tag{3.11}$$

for every $(n_1,n_2) \in \mathbb{Z}^2$. With this convention the data $f[t_1,t_2]$ itself can be viewed as samples $f(\frac{t_1}{n},\frac{t_2}{n})$ of f, the inverse (continuous) Fourier transform of $\hat{f}$ restricted to the fundamental cell.

Using this continuous representation of the data, along with equation (3.12) in the case

when the modulo is triggered in equation (3.11), $c^D(j,\ell,k)$ obeys

$$c^D(j,\ell,k) = \frac{1}{n^2} \sum_{n_1,n_2 \text{ in tile}} e^{i2\pi(k_1 n_1/L_{1,j} + k_2 n_2/L_{2,j})}\, \hat{\varphi}^D_{j,\ell}(2\pi n_1, 2\pi n_2)\, \hat{f}(2\pi n_1, 2\pi n_2)$$

and since $\hat{\varphi}_{j,\ell}$ is compactly supported, one can extend the sum above to $(n_1,n_2) \in \mathbb{Z}^2$. Introduce the Dirac comb

$$c(\omega_1,\omega_2) = \sum_{n_1\in\mathbb{Z}} \sum_{n_2\in\mathbb{Z}} \delta(\omega_1 - 2\pi n_1)\delta(\omega_2 - 2\pi n_2).$$

and rewrite $c^D(j,\ell,k)$ as

$$c^D(j,\ell,k) = \frac{1}{n^2} \int_{\mathbb{R}^2} e^{i\omega_1 \frac{k_1}{L_{1,j}}}\, e^{i\omega_2 \frac{k_2}{L_{2,j}}}\, c(\omega)\hat{\varphi}^D_{j,\ell}(\omega)\hat{f}(\omega)\, d\omega.$$

Our claim follows from Parseval's identity which states that $\int \hat{u}\overline{\hat{v}} = (2\pi)^2 \int u\overline{v}$. Indeed, the inverse Fourier transform of $\hat{f}$ is given by

$$\mathcal{F}^{-1}(\hat{f}(\omega))(x) = \sum_{t_1=0}^{n-1} \sum_{t_2=0}^{n-1} \delta(x_1 - \frac{t_1}{n})\delta(x_2 - \frac{t_2}{n}) f[t_1,t_2],$$

while for the other

$$\mathcal{F}^{-1}(e^{-i\omega_1 \frac{k_1}{L_{1,j}}}\, e^{-i\omega_2 \frac{k_2}{L_{2,j}}}\, c(\omega)\hat{\varphi}^D_{j,\ell}(\omega))(x) = \frac{1}{(2\pi)^2}\varphi^\sharp_{j,\ell}(x_1 - \frac{k_1}{L_{1,j}}, x_2 - \frac{k_2}{L_{2,j}}).$$

The Parseval formula then gives (3.10). For the North and South quadrants, the proof is identical after swapping $L_{1,j}$ and $L_{2,j}$.

□

Notice that the actual value of x_μ, the center of $\varphi_\mu(x)$ in physical space, is implicit in formula (3.10). If φ_μ is centered at the origin when $k_1 = k_2 = 0$, then

$$x_\mu = (\frac{k_1}{L_{1,j}}, \frac{k_2}{L_{2,j}})$$

when the angle is $-\pi/4 \le \theta_\ell < \pi/4$, and

$$x_\mu = (\frac{k_1}{L_{2,j}}, \frac{k_2}{L_{1,j}})$$

for angles $\pi/4 \leq \theta_\ell < 3\pi/4$.

3.2.2 Isometry and Inversion

In practice the curvelet coefficients are normalized as follows,

$$c^{D,N}(j,\ell,k) = \frac{n}{\sqrt{L_{1,j}L_{2,j}}} c^D(j,\ell,k),$$

where $L_{1,j}, L_{2,j}$ are the sidelengths of the parallelogram $\mathcal{P}_{j,\ell}$. Equipped with this normalization, we have the Plancherel relation

$$\sum_{t_1,t_2} |f[t_1,t_2]|^2 = \sum_{j,\ell,k} |c^{D,N}(j,\ell,k)|^2.$$

This is easily proved by noticing that every step of the transform is isometric.

- The discrete Fourier transform, properly normalized,

  $$f[t_1,t_2] \to \frac{1}{n}\hat{f}[n_1,n_2]$$

  is an isometry (and unitary).

- The decomposition into different scale-angle subbands,

  $$\hat{f}[n_1,n_2] \to \{\tilde{U}_{j,\ell}[n_1,n_2]\hat{f}[n_1,n_2]\}_{j,\ell}$$

  is an isometry because the windows $\tilde{U}_{j,\ell}$ are constructed to obey $\sum_{j=0}^{J}\sum_\ell \tilde{U}_{j,\ell}(\omega)^2 = 1$.

- The wrapping transformation is only a relabeling of the frequency samples, thereby, preserving ℓ^2 norms.

- The local inverse Fourier transform (3.9) is an isometry when properly normalized by $\frac{1}{\sqrt{L_{1,j}L_{2,j}}}$.

Owing to this isometry property, the inverse curvelet transform is simply computed as the adjoint of the forward transform. Adjoints can typically be computed by "reversing" all the operations of the direct transform. In our case,

1. For each scale/angle pair (j, ℓ), perform a (properly normalized) 2-D FFT of each array $c^{D,N}(j, \ell, k)$, and obtain $W(\tilde{U}_{j,\ell}\hat{f})[n_1, n_2]$.

2. For each scale/angle pair (j, ℓ), multiply the array $W(\tilde{U}_{j,\ell}\hat{f})[n_1, n_2]$ by the corresponding wrapped curvelet $W(\tilde{U}_{j,\ell})[n_1, n_2]$ which gives

$$W(|\tilde{U}_{j,\ell}|^2\hat{f})[n_1, n_2].$$

3. Unwrap each array $W(|\tilde{U}_{j,\ell}|^2\hat{f})[n_1, n_2]$ on the frequency grid and add them all together. This recovers $\hat{f}[n_1, n_2]$.

4. Finally, take a 2-D inverse FFT to get $f[t_1, t_2]$.

In the wrapping approach, both the forward and inverse transform are computed in $O(n^2 \log n)$ operations, and require $O(n^2)$ storage.

3.3 Extensions

3.3.1 Curvelets at the Finest Scale

The design of appropriate basis functions at the finest scale, or outermost dyadic corona, is not as straightforward for directional transforms like curvelets as it is for 1-D or 2-D tensor-based wavelets. This is a *sampling* issue. If a fine-scale curvelet is sampled too coarsely, the pixelization will make it look like a checkerboard and it will not be clear in which direction it oscillates anymore. In the frequency domain, the wedge-shaped support does not fit in the fundamental cell and its periodization introduces energy at unwanted angles.

The problem can be solved by assigning wavelets to the finest level. When $j = J$, the unique sampled window $\tilde{U}_J[n_1, n_2]$ is so constructed that its square forms a partition of unity, together with the curvelet windows. A full 2D inverse FFT can then be performed to obtain the wavelet coefficients. This highpass filtering is very simple but goes against the philosophy of directional basis elements at fine scale. Wavelets at the finest scale are illustrated in Figure 3.7 (top row).

In this section, we present the next simplest solution to the design of faithful curvelets at the finest scale. For simplicity let us adopt the sampling scheme of the wrapping implementation, but a parallel discussion can be made for the USFFT-based transform. As above,

denote by J the finest level. By construction, the standard curvelet window $\tilde{U}_{j,\ell}[n_1, n_2]$ is obtained by sampling a continuous profile $\tilde{U}_{j,\ell}(\omega_1, \omega_2)$ at $\omega_1 = 2\pi n_1$, $\omega_2 = 2\pi n_2$. When $j = J$, the profile $\tilde{U}_{j,\ell}$ overlaps the border of the fundamental cell but can still be sampled according to the formula

$$\tilde{U}_{J,\ell}[(n_1 + \frac{n}{2}) \bmod n - \frac{n}{2}, (n_2 + \frac{n}{2}) \bmod n - \frac{n}{2}] = \tilde{U}_{J,\ell}(2\pi n_1, 2\pi n_2). \tag{3.12}$$

The indices n_1, n_2 are still chosen such that $\tilde{U}_{J,\ell}$ is evaluated on its support. The latter is by construction sufficiently small so that no confusion occurs when taking modulos. In effect we have just copied $\tilde{U}_{J,\ell}$ by periodicity inside the fundamental cell. The windows $\tilde{U}_{J,\ell}(\omega_1, \omega_2)$ must be chosen adequately so that the discrete arrays $\tilde{U}_{J,\ell}[n_1, n_2]$, now with $n_1, n_2 = -n/2 \ldots n/2 - 1$, obey the isometry property together with the other windows,

$$\sum_{j=0}^{J} \sum_{\ell} |\tilde{U}_{j,\ell}[n_1, n_2]|^2 = 1.$$

In fact, this is the case if $\tilde{U}_{J,\ell}$ is chosen as in Section 3.1 (after an appropriate rescaling).

Periodization in frequency amounts to sampling in space, so finest-scale curvelets are just undersampled standard curvelets. This is illustrated in Figure 3.7 (middle row). What do we loose in terms of aliasing? Spilling over by periodicity is inevitable, but here the aliased tail consists of essentially only one-third of the frequency support. Observe in Figure 3.7 (middle right) that a large fraction of the energy of the discrete curvelet still lies within the fundamental cell. Numerically, the non-aliased part amounts to about 92.4% of the total squared ℓ^2 norm $\|\varphi^D_{j,\ell,k}\|^2_{\ell^2}$. The "checkerboard" look of undersampled curvelets, mentioned above, is shown in Figure 3.7 (bottom right).

Accordingly, the definition of wrapping of an array $d[n_1, n_2]$, in the presence of undersampled curvelets, is modified to read:

$$Wd[n_1 \bmod L_{1,j}, n_2 \bmod L_{2,j}] = d[(n_1 + \frac{n}{2}) \bmod n - \frac{n}{2}, (n_2 + \frac{n}{2}) \bmod n - \frac{n}{2}] \tag{3.13}$$

The new modulo that appears in the above equation (compare with (3.7)) prevents data queries outside $[0, n]^2$, which would otherwise happen if equation (3.7) were used naively. Instead, data is folded back by periodicity onto the fundamental cell, ultimately resulting

in aliased basis functions.

The definitions of forward and inverse curvelet transforms, as well as their properties, otherwise go unchanged. Proposition 3.1 and its proof do not have to be changed either: they are already compatible with equation (3.13).

3.3.2 Windows over Junctions Between Quadrants

The construction of windows $\tilde{U}_{j,\ell}$ explained in Section 3.1.1 make up an orthonormal partitioning of unity as long as the window is supported near wedges that do not touch neither of the two diagonals. There are 8 "corner" wedges per scale calling for a special treatment, and corresponding to angles near $\pm\pi/4$ and $\pm 3\pi/4$, see Figure 3.3 on the left. In these exceptional cases, creating a partition of unity is not as straightforward. This is the topic of this section.

It is best to follow Figure 3.3 while reading this paragraph. Consider a trapezoid in the top quadrant and corresponding to an angle near $3\pi/4$ as in the figure. The grey trapezoid is the corner wedge near which the curvelet is supported, but the actual support of the curvelet is the nonconvex hexagon bounded by the dash-dotted line. As before, the corner curvelet window is given as a product of the radial window W_j and of the angular window $V_{j,\ell}$,

$$\hat{\varphi}^D_{j,\ell}(\omega) = W_j(\omega)V_{j,\ell}(\omega).$$

We decompose the corner window $V_{j,\ell}$ into a left-half and a right-half. The right-half is given by the standard construction presented earlier. It is a function of $\frac{\omega_1}{\omega_2}$. The left-half of the window is constructed as a member of a square-root of a partition of unity designed in a frame rotated by 45 degrees with respect to the Cartesian axes. The left-half of the window is a function of $\frac{\omega_1+\omega_2}{\omega_1-\omega_2}$. The left and right-halves agree on the line where they are stitched together (on the figure, it is the tilted line, first to the right of the diagonal $\omega_1 = -\omega_2$). Along the border line, they are both equal to one and they have at least a couple of vanishing derivatives in all directions. Again, the partition of unity can be designed so that all these derivatives are zero. By construction, our set of windows obeys the partition of unity property, equation (3.2).

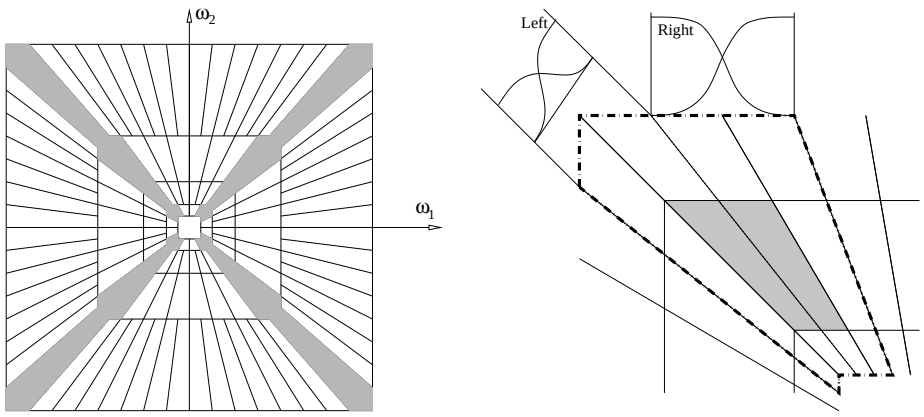

Figure 3.3: Left: The corner wedges appear in grey. Right: Detail of the construction of a partition of unity over the junction between quadrants.

3.3.3 Other Frequency Tilings

The construction of curvelets is based on a polar dyadic-parabolic partition of the frequency plane, also called FIO tiling, as explained in Section 2.2. However, the approach is flexible, and can be used with a variety of choices of parallelepipedal tilings, for example, including based on principles besides parabolic scaling. For example:

- A *directional wavelet* transform is obtained if, instead of dividing each dyadic corona into $C \cdot 2^{\lfloor j/2 \rfloor}$ angles, we divide it into a constant number, say 8 or 16 angles, regardless of scale as in [71]. This can be realized by dropping the requirement that wedges be split as scale increases.

- A *ridgelet* transform is obtained by subdividing each dyadic corona into $C \cdot 2^j$ angles. This can be achieved by subdividing every angular wedge every time the scale index j increases (not just every other time, as for curvelets.)

- A *Gabor* analysis is obtained if, instead of considering bandpass concentric annuli of thickness increasing like a power of two, we consider the thickness to be the same for all annuli. In other words, coronae with fixed width are substituted for dyadic coronae. The number of wedges into which an annulus should be divided is proportional to its length, or equivalently, its distance to the origin.

- More generally, one can create an adaptive partitioning of the frequency plane that best matches the features of the analyzed image. This is the construction of *ridgelet*

packets as explained in [40]. A best basis strategy can then be overlaid on the packet construction to find the optimal partitioning in the sense that it minimizes an additive measure of "entropy", or sparsity.

In all these cases the wrapping strategy carries over without essential modifications and yield tight or nearly tight frames. The design problem is reduced to the construction of a smooth partition of unity that indicates the desired frequency tiling.

3.3.4 Higher Dimensions

Curvelets exist in any dimension [15]. In 3 dimensions for example, curvelets are little plates of sidelength about $2^{-j/2}$ in two directions and thickness about 2^{-j} in the orthonormal direction. They vary smoothly in the two long directions and oscillate in the short one (the 3D parabolic scaling matrix is of the form $\mathrm{diag}(2^{-j/2}, 2^{-j/2}, 2^{-j})$). Just as 2D curvelets provide optimally efficient representations of 2D objects with singularities along smooth curves, 3D curvelets would provide efficient representations of 3D objects with singularities along smooth 2D surfaces, and more generally, of objects with singularities along smooth manifolds of codimension 1 in higher dimensions.

The algorithms for 3D discrete curvelet transforms are similar to their 2D analogs. We first decompose the object into dyadic annuli based on concentric cubes. Each annulus is subdivided into trapezoidal regions obeying the usual frequency parabolic scaling (one long and two short directions), see Figure 3.4. (Note that they are now 6 components corresponding to the 6 faces of the cube.)

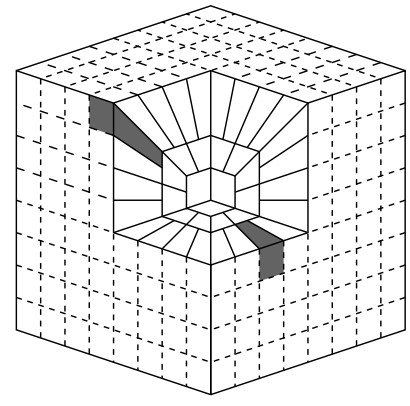

Figure 3.4: The dyadic-parabolic frequency tiling in 3D. Curvelets are supported near the gray regions.

Both transforms carry over to 3 dimensions and we only rehearse the minor modifications.

1. The 3D FDCT via wrapping just wraps the 3D parallelepipeds instead of their 2D analogs.

2. The construction of junction windows (described in Section 3.3.2 for 2D FDCTs) is a little more delicate since one needs to consider more cases. One possible solution is to develop a partition of unity over the unit sphere which is then mapped onto the cube. The detailed algorithm and numerical results of the 3D transform will be presented in a future report.

In short, 3D FDCTs follow exactly the same architecture as 2D FDCTs, and the forward, adjoint, and inverse transforms all run in $O(N \log N)$ for Cartesian arrays of size $N = n^3$ voxels.

3.3.5 Nonperiodic Image Boundaries

An (unfortunate) consequence of using the DFT to define our transform is that the image is implicitly considered as a periodic array. The leftmost and rightmost pixels in a given row, or the top and bottom pixels in a given column, are considered immediate neighbors as much as ordinary adjacent pixels are. By construction, a substantial number of basis functions appear to be supported on two (or more) very distant regions of the image, because they overlap the image boundary and get copied by periodicity. Let us call them "boundary curvelets."

Periodization may result in unwanted curvelet-looking artifacts near the image boundary, for example in image denoising experiments. The reason for the presence of these artifacts, however, is not the same for curvelets and for wavelets. In order to understand this phenomenon, we need to sort curvelets according to their orientation.

1. Boundary curvelets that are *aligned* with a boundary edge mostly respond to the artificial discontinuity created by periodization. Since the basis elements very closely follow the boundary, the visual effect of a big coefficient is minor.

2. Boundary curvelets *misaligned* with respect to the boundary edge are assigned big coefficients when they respond to geometrical structure on the *opposite side* of the

image, across the edge. This causes the most severe visual artifacts.

In the remainder of this section, we present a few (somewhat naive) solutions to artifacts of type 2, when boundary curvelets are misaligned.

The most obvious remedy is to pad the image with zeros to make it twice larger in both directions. The curvelet transform is then applied to the extended image, increasing the redundancy by a factor 4. The blank surrounding region is large enough to prevent boundary curvelets from wrapping around. The inverse or adjoint transform would then would have an extra step, clipping off the extra pixels.

If we postulate that artifacts of type 2 are caused by boundary curvelets forming an angle greater than 45 degrees with the edge, then it is not necessary to zeropad in all directions. The image should only be extended horizontally for mostly horizontal curvelets, and vertically for mostly vertical curvelets. The zeropadding will make the image twice larger in only one direction, depending on the orientation of the subband considered. In this case, the increase in redundancy is only of a factor 2.

In principle it would be advantageous to make the width of the zeropadding not only angle-dependent, but also scale-dependent. More precisely, the width of the padding does not have to be bigger than a factor times the length of misaligned curvelets, i.e., $C \cdot 2^{-\lfloor j/2 \rfloor}$. The gain in redundancy would be obvious. There is a complication, however, in considering scale-dependent or even angle-dependent paddings. Different subbands will correspond to different grids and extra care will be needed to properly re-design the transform to make it an isometry. It will be necessary to rethink the notion of discrete partition of unity to accommodate interpolation between different grids.

We have not pursued this issue much further, but a better handling of image boundaries would improve the current architecture of the curvelet transform for image processing applications.

3.4 Numerical Examples

We start this section by displaying a few curvelets in both the spatial and the frequency domain, see Figures 3.5 (coarsest scale curvelets), 3.6 and 3.7 (curvelets at the finest level where one can choose between wavelets and curvelets). Localization in both space and frequency is apparent. The digital curvelets seem faithful to their continuous analog. In the

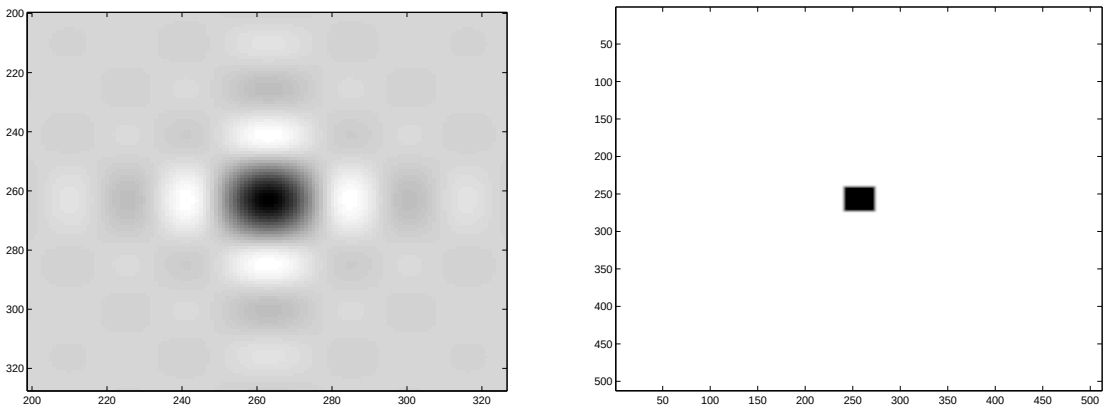

(a) (b)

Figure 3.5: At the coarsest level, curvelets are nondirectional and are Meyer scaling functions. (a) Spatial-side. The color map is as follows: white is most negative, zero corresponds to some tone of grey, and black is most positive. (b) Frequency-side (modulus of the Fourier transform). The level of grey indicates values from zero (white) to one (black).

spatial domain, they are smooth along and oscillatory across the ridge. In the frequency domain, they are sharply localized.

Next, Tables 3.1 and 3.2 report the running time of both FDCTs on a sequence of arrays of increasing size. T_{Fwd}, T_{Inv} and T_{Adj} are running times of the forward, inverse and adjoint transforms respectively (we only give T_{Inv} for the FDCT via wrapping since the inverse is the same as the adjoint). The column T_{Fwd}/T_{FFT} gives the ratio between the running time of the FDCT and that of the FFT on an array of the same size. The accuracy or ℓ^2-error is computed as $\|f - C_{Inv}C_{Fwd}f\|_{\ell^2}/\|f\|_{\ell^2}$ where C_{Inv} and C_{Fwd} are the the forward and inverse FDCTs. The FDCT via wrapping achieves machine accuracy because of the exact numerical tightness of the digital transform. The FDCT via USFFT also achieves high accuracy, i.e. of the order of 10^{-6}. Although both transforms have low running times, the USFFT transform is somewhat slower; this is due to the interpolation step in the forward transform and to the CG iterations in the inverse transform.

Image size	T_{Fwd}(s)	T_{Inv}(s)	T_{Fwd}/T_{FFT}	ℓ^2 error
128×128	0.040458	0.039520	11.2383	4.5450e-16
256×256	0.174807	0.176519	8.8286	4.8230e-16
512×512	0.829820	0.868141	6.0793	4.8908e-16
1024×1024	4.394066	4.482452	7.7224	5.6303e-16
2048×2048	20.01692	23.02144	7.7567	6.3018e-16

Table 3.1: Running time and error for the wrapping-based transform.

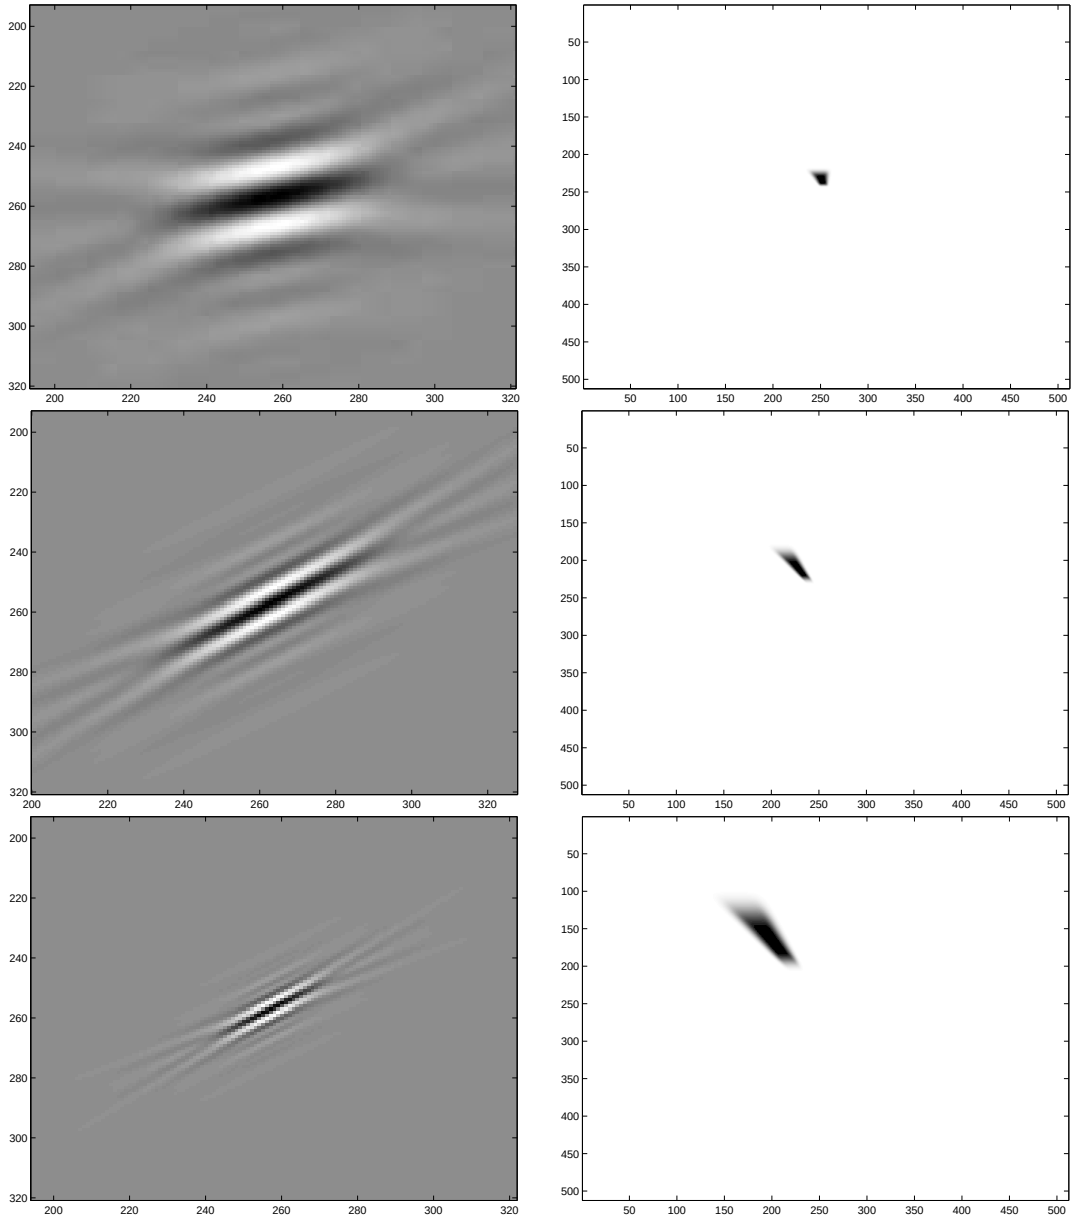

Figure 3.6: Curvelets at increasingly fine scales. The left panels represent curvelets (real part) in the spatial domain (as functions of the spatial variable x). The right panels show the modulus of the Fourier transform (as functions of the frequency variable ω). The color map is the same as in Figure 3.5.

We then illustrate the potential of FDCTs with several examples. The wrapping-based implementation has been used for all experiments. In the first example, we compare the decay of the coefficients of the curvelet and various wavelet representations on images with curve-like singularities. Our first input image—shown in Figure 3.8 (a)—is singular along a smooth curve and is otherwise perfectly smooth (this image is de-aliased to remove the

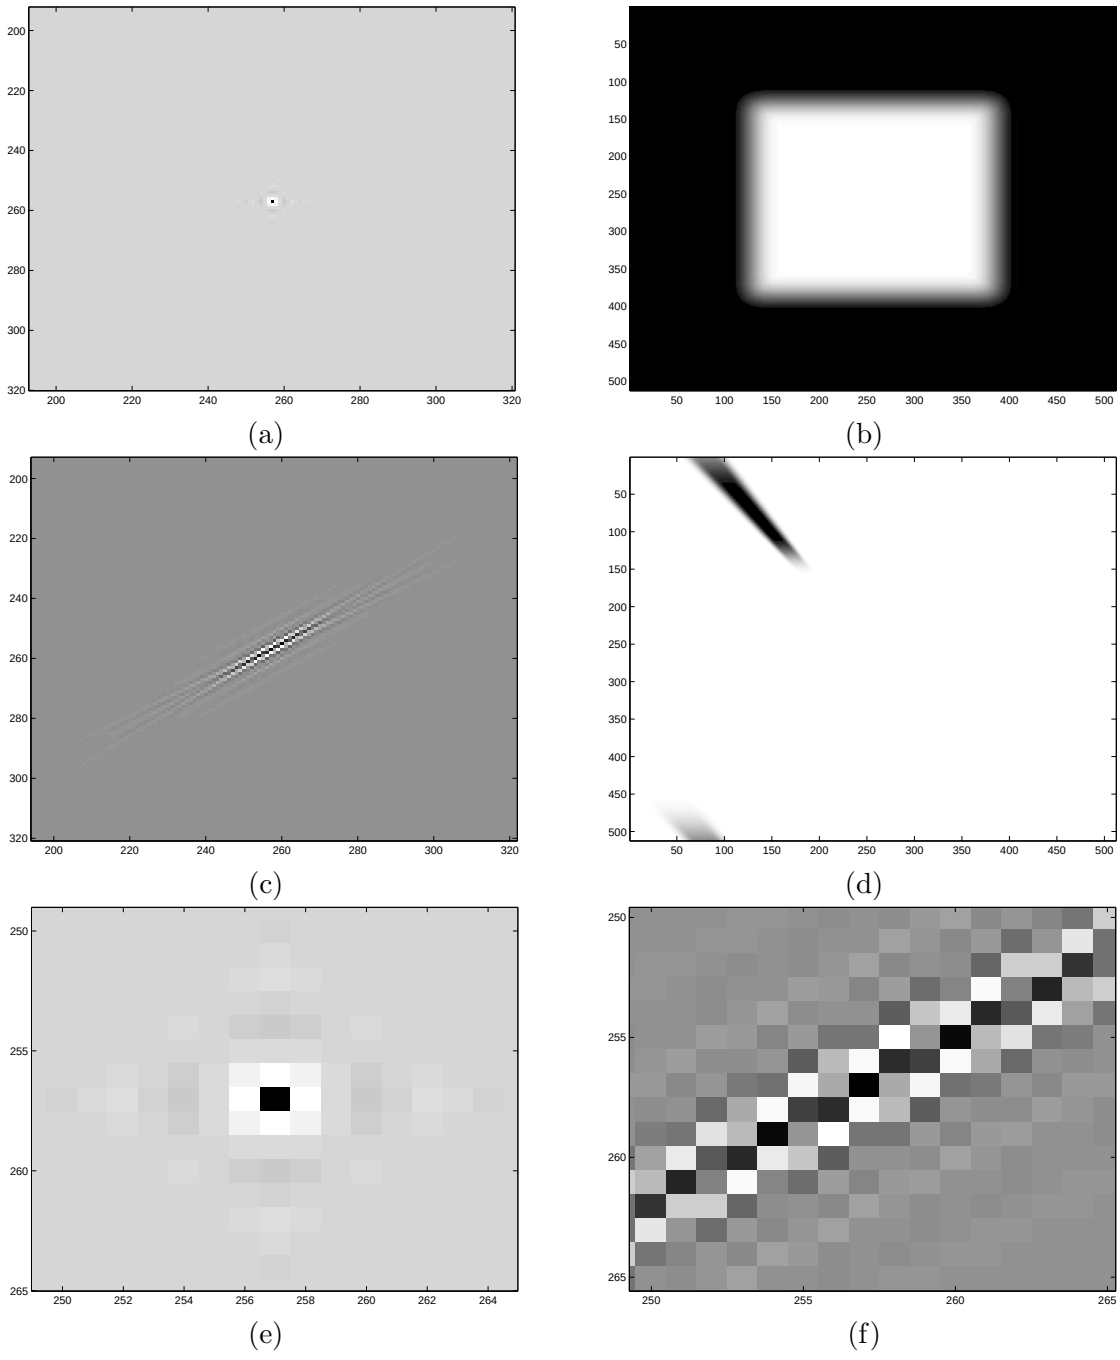

Figure 3.7: Wavelets and curvelets at the finest scale. Meyer wavelet in space (a) and frequency (b). Undersampled curvelet in space (c) and frequency (d). (e) Zoom of (a). (f) Zoom of (c).

artifacts due to pixelization). To compensate for the redundancy of the curvelet transform and to display a meaningful comparison, we extract a fraction of the entries of the curvelet coefficient table so that the number of curvelet and wavelet coefficients is identical. The extracted curvelet entries are renormalized to preserve the overall ℓ^2 norm. Figure 3.8 (b)

Image size	T_{Fwd}(s)	T_{Adj}(s)	T_{Inv}(s)	T_{Fwd}/T_{FFT}	ℓ^2 error
128×128	0.088832	0.091578	1.006522	24.6756	1.4430e-06
256×256	0.376838	0.390533	4.002353	19.0322	8.8154e-07
512×512	2.487052	2.579102	35.09599	18.2202	5.3195e-07
1024×1024	16.47702	16.87764	129.3631	28.9579	3.2390e-07
2048×2048	62.42980	65.09365	566.1732	24.1920	3.4305e-06

Table 3.2: Running time and error for the USFFT-based transform.

shows the values of the coefficients sorted in decreasing order of magnitude. The faster the decay, the better. The sparsity analysis is complemented by the quantitative study of partial reconstructions of f, where we have again used redundancy compensation as explained above. Figure 3.8 (c) shows the PSNR of best m-term approximation,

$$PSNR = 20 \log_{10} \left(\frac{\max(f(x)) - \min(f(x))}{\|f - f_m\|_2} \right) \qquad (dB)$$

where f_m is the partial reconstruction of f using the m largest coefficients in magnitude, in the curvelet (or wavelet) expansion (note that because of the redundancy of the FDCT, there are better ways of obtaining partial reconstructions).

The second input image—shown in Figure 3.9 (a)—is a synthetic seismogram corresponding to the acoustic response of a one-dimensional layered medium to a point source. The decay of the coefficients and the partial reconstruction error for this image are shown in Figure 3.9 (b) and (c) respectively. Our experiments suggest that FDCTs outperform, by a significant margin, traditional wavelet representations on these types of image data. Synthetic seismic images seem to be the ideal setting for curvelets because they are prepared as solutions to a wave equation in simple layered media, with a bandlimited point excitation. The solution itself is therefore very close to being bandlimited. We are in the setting of proposition 3.1: when the data are oscillatory yet properly sampled, curvelets are expected to be completely faithful to the continuous transform, explaining the good denoising performance.

The second example is denoising. The original image is the seismogram used in the previous example (see Figure 3.9 (a)). The noise-to-signal ratio is set to 10%, which corresponds to PSNR = 20.0 dB. A denoising algorithm based on our curvelet transform results in an image with PSNR = 37.6 dB. (see Figure 3.10 (c)) while a traditional wavelet denosing algorithm (Symmlet 8 in WaveLab, shift-invariant hard thresholding at 2.5σ) gives

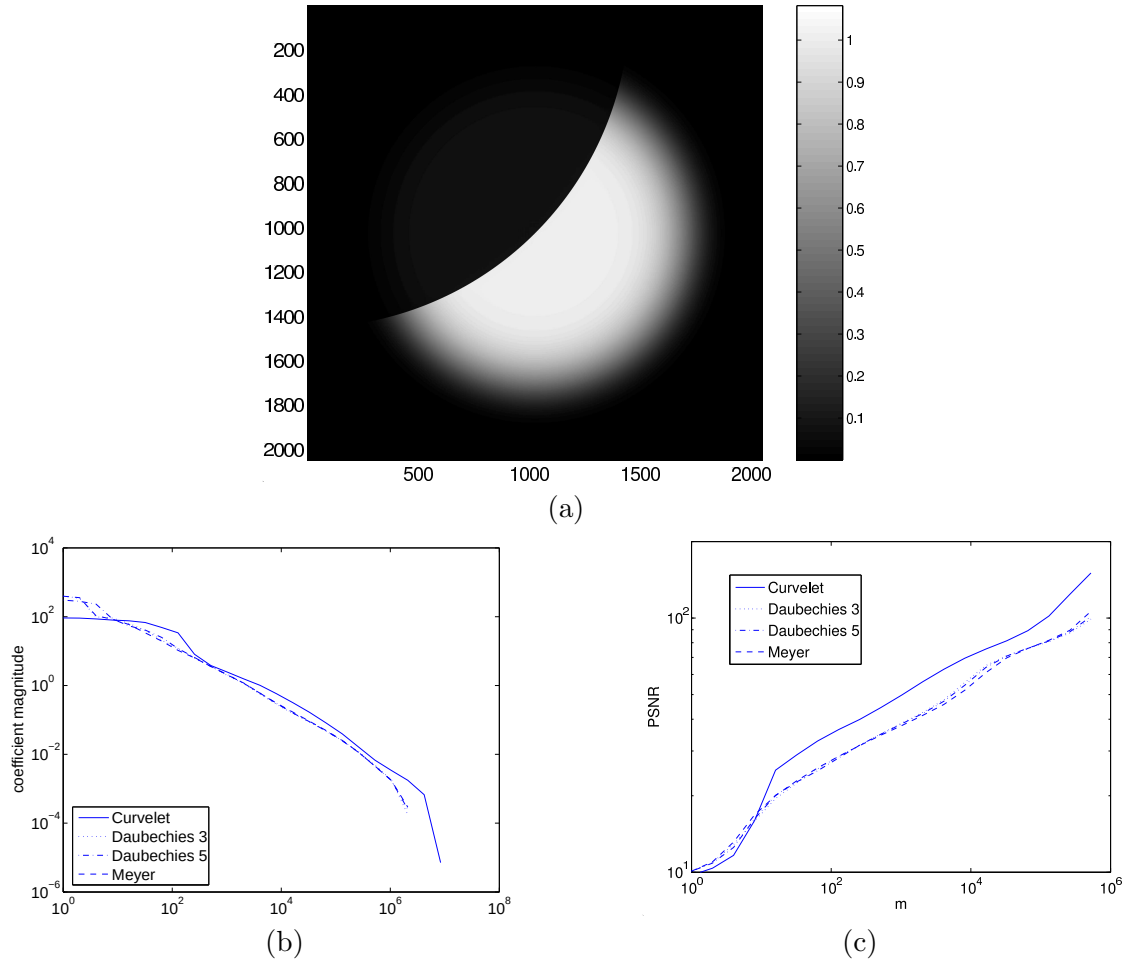

Figure 3.8: Sparsity analysis of the curvelet and wavelet representations of a singular object. (a) Input image. (b) Magnitude of the coefficients sorted in descending order. (c) PSNR for partial reconstruction with the m largest coefficients in magnitude. The horizontal line at 40 dB indicates a typical "visually acceptable" level of reconstruction.

PSNR = 30.8 dB. (see Figure 3.10 (d)). The curvelet denoising algorithm used above is a simple shift-invariant block-thresholding of the wrapping-based curvelet transform (with curvelets at the finest scale) and is available as Matlab code in CurveLab. (For an image of size 1024×512, the whole procedure runs in less than 90 seconds on a standard desktop.)

In the introduction chapter, we pointed out that curvelets were especially well-adapted to simultaneously represent the solution operators to large classes of wave equations and the wavefields that are solutions to those equations. In our third example, we consider the constant coefficient second-order wave equation with periodic boundary condition

$$u_{tt} - \Delta u = 0 \quad x \in [0,1) \times [0,1).$$

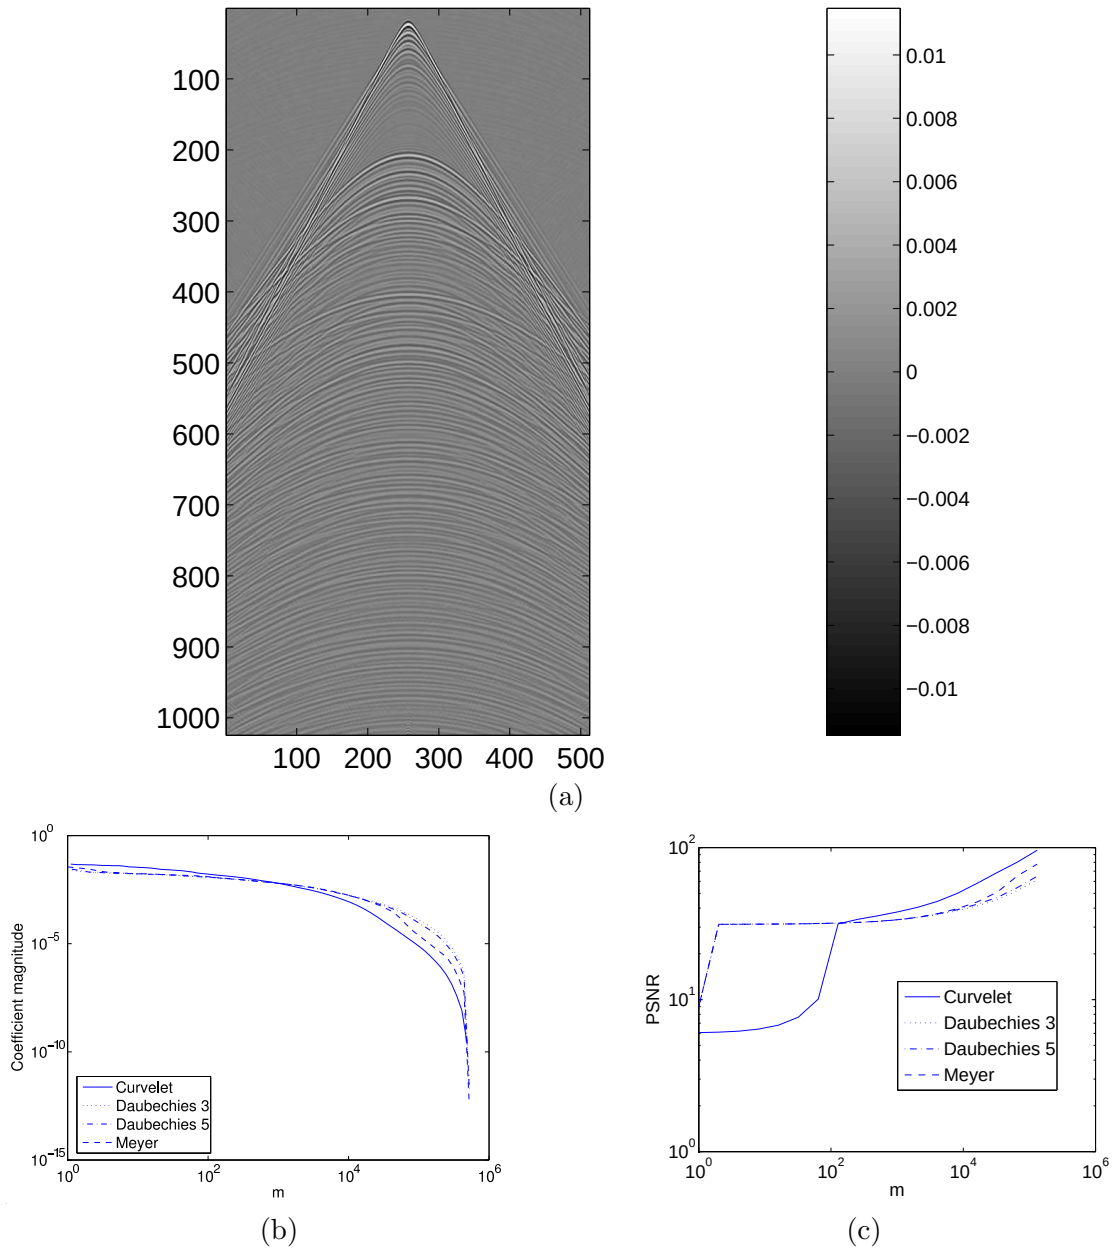

Figure 3.9: Sparsity analysis of the curvelet and wavelet representations of a seismogram. (a) Synthetic seismogram corresponding to the acoustic response of a one-dimensional layered medium to a point source, courtesy of Eric Verschuur and Felix Herrmann. The x-axis is the offset from the source and the y-axis is time. (b) Decay of the coefficients. (c) Partial reconstruction error, measured in PSNR.

We discretize the domain with a 512-by-512 Cartesian grid, and take as initial wavefield a delta function located at the center of the domain, see Figure 3.11 (a). The solution at a later time is known analytically, and may therefore be computed exactly. We use the FDCT to compress the wavefield at time $t = 0.25$ and $t = 0.75$. Figures 3.11 (b) and (c) show the

approximate wavefields reconstructed from only 1.25% of the curvelet coefficients. In both cases, the relative ℓ^2 error is about 10^{-5}.

We have seen that the wavefield is well-approximated by just a few curvelets and now study the compressibility of the wave propagator E_t. For simplicity, assume E_t acts on scalar wavefields. From a theoretical point of view, it is known that the entries of $E_t(\mu, \mu') = \langle \varphi_\mu, E_t \varphi_{\mu'} \rangle$ taken from an arbitrary row (fixed μ) or column (fixed μ') decay faster than any negative power law. Figure 3.11 (d) plots the decay of the matrix coefficients (sorted by decreasing magnitude) for several columns of the propagator matrix E_t at $t = 0.75$ while (e) plots the relative truncation error for those same columns. "Scale" in the legend refers to the scale j' corresponding to μ', the index of the column. Observe that for every column, we achieve a relative error of order 10^{-5} by using about 1% of the largest curvelet coefficients. The data are shown as is; no compensation for redundandy has been made in this experiment.

3.5 Discussion

The transform introduced in this chapter was designed with the goal of being as faithful to continuous curvelets as possible. The main step of the transform is to window the data in frequency with prescribed windows, sampled on the same grid as the data. This sampling in frequency is the *only* distortion that curvelets incur in the digital transforms. This issue is inevitable but minor, since it is equivalent to periodization in space where curvelets decay fast. Recall that the other potential source of error, spatial sampling, is a nonissue here since curvelets are nearly bandlimited.

The wrapping variant is to our knowledge the fastest curvelet transform currently available. Computing a direct or inverse transform in C++ takes about the same time as 6 to 10 FFTs using FFTW (available at http://www.fftw.org), which can hardly be improved upon.

3.5.1 Relationships with Other Works

The notion of directional multiscale transform originated independently in different fields in the early nineties. Without the claim of being exhaustive, let us only mention continuous wavelet theory [67] and steerable pyramids in the field of computer vision [71, 70]. The latter

approach was the first practical, data-friendly strategy to extract information at different scales and angles.

A more recent, very interesting attempt at implementing low-redundancy curvelets, was introduced by Minh Do and Martin Vetterli, in [32]. The construction is based on a filterbank decomposition of the image in both scale and angle. The resulting basis functions are called "contourlets," and form a tight frame with redundancy 4/3. The contourlet transform has a very fast $O(n^2 \log n)$ implementation as well, at least when contourlets are selected to be compactly supported. The only problem with this construction is that it is not faithful to the idea of the curvelet transform in the sense that for most choices of filters in the angular filterbank, contourlets are not sharply localized in frequency. On the practical side, this means that contourlets lack smoothness along the ridge in the spatial domain and exhibit spurious oscillations which may be of source of numerous problems, especially if one wants to use these transforms for scientific computing. On the theoretical side and to the best of our knowledge, contourlets do not allow to formulate as strong theorems in approximation and operator theory as in [15, 20].

The idea of using concentric squares and shears is also central to the construction of tight-frames of "shearlets", by Guo, Kutyniok, Labate, Lim, Weiss and Wilson in a recent series of papers [44, 45, 56] starting with [44]. In these papers, they show how to built wavelets or multiwavelets from *composite dilations* and translations. The architecture is very similar to that of curvelets.

3.5.2 Possible Applications

Just as the wavelet transform has been deployed a countless number of times in many fields of science and technology, we expect fast digital curvelet transforms to be widely applicable—especially in the field of image processing and scientific computing.

In image analysis for example, the curvelet transform may be used for the compression of image data, for the enhancement and restoration of images as acquired by many common data acquisition devices (e.g., CT scanners), and for postprocessing applications such as extracting patterns from large digital images, detecting features embedded in very noisy images, enhancing low contrast images, or registering a series of images acquired with very different types of sensors.

Curvelet-based seismic imaging already is already a very active field of research, see for

example the recent papers [47, 49] as well as several expanded abstracts by Felix Herrmann and his collaborators, currently available at http://slim.eos.ubc.ca/.

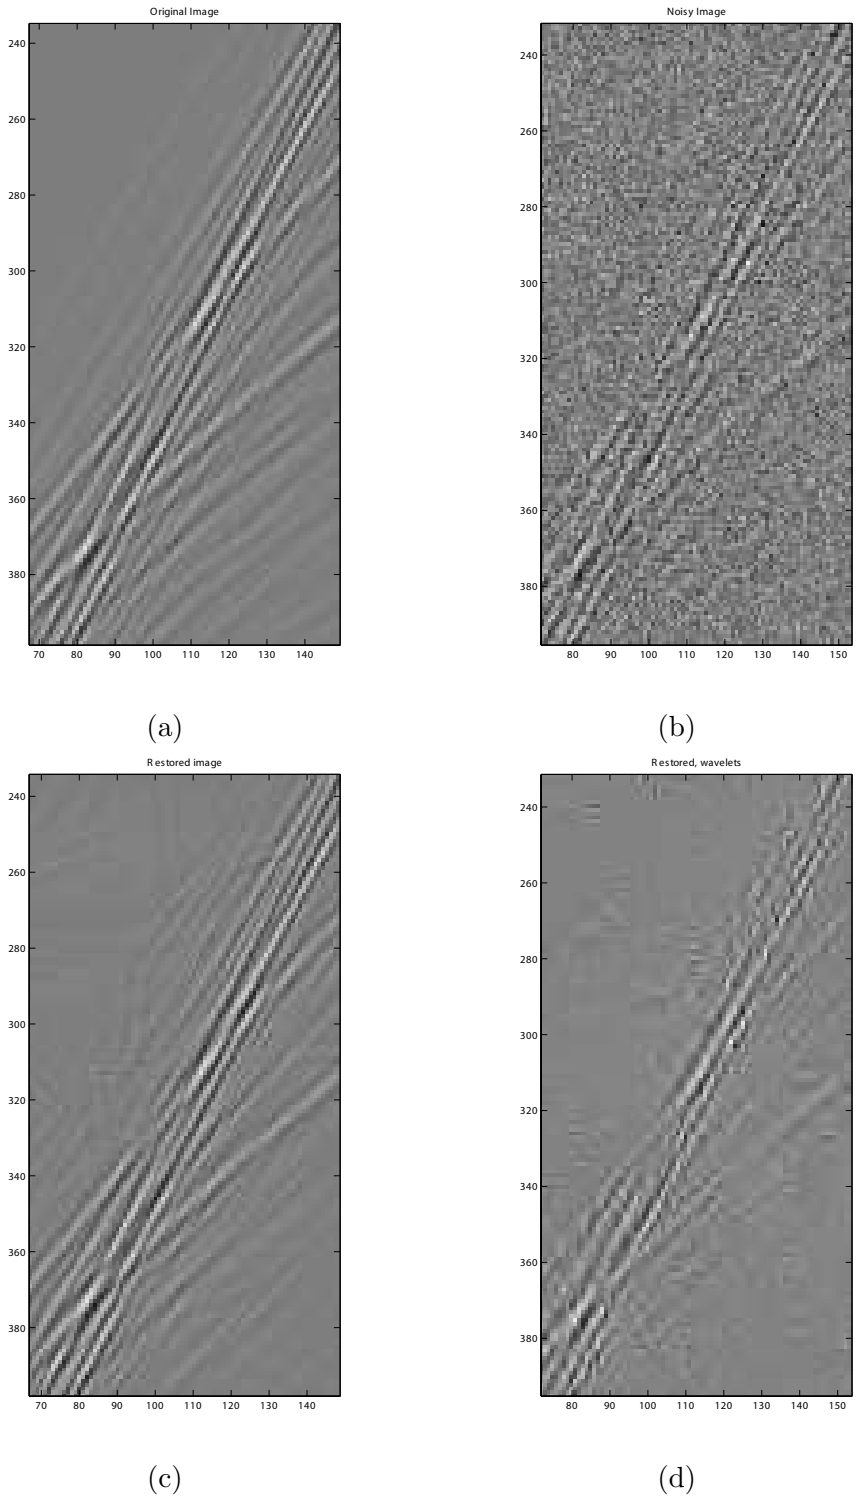

Figure 3.10: Image denoising using curvelets. (a) The Original image (zoom). (b) Noisy image (Gaussian white noise with $\sigma = 10\%$ of the maximum intensity), PSNR = 20.0 dB. (c) Denoised image using curvelets, PSNR = 37.6 dB. (d) Denoised image using wavelets, PSNR = 30.8 dB.

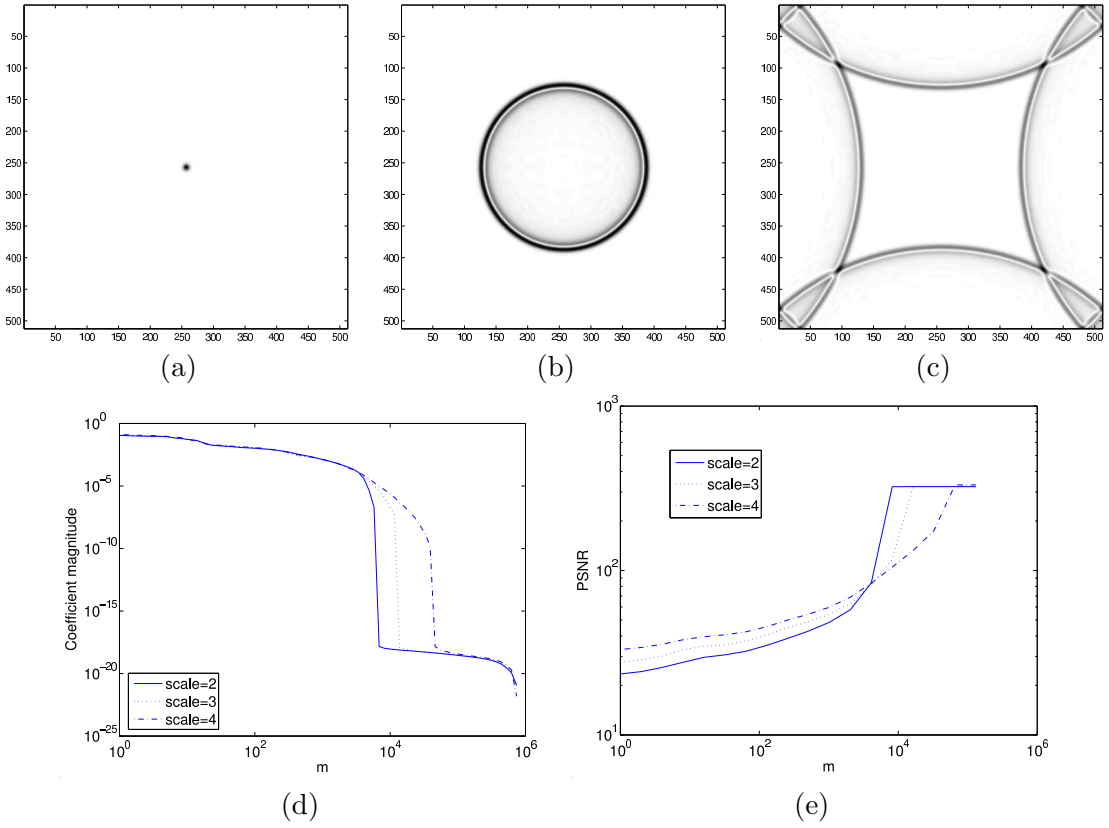

Figure 3.11: Compression of the wavefield and of the solution operator to the wave equation with periodic boundary conditions. (a) The initial condition is a delta function located at the center of the domain. (b) Approximate solution at $t = 0.25$. (c) Approximate solution at $t = 0.75$. Both approximations only use 1.25% of nonzero curvelet coefficients. (d) Magnitude of the matrix entries (rearranged in descending order) of the solution operator E_t at $t = 0.75$ taken from three columns corresponding to three curvelets at various scales. (e) For the same three columns, truncation error obtained by keeping the m largest entries, measured in PSNR.

Chapter 4

Wave atoms and time upscaling of wave equations

In this chapter, we show how wave atoms can be used to formulate a fast algorithm for wave propagation in smooth, periodic, 2D inhomogeneous media.

We already know that wave atoms offer a uniquely structured representation of the time-dependent Green's function in the sense that the resulting matrix is *universally sparse* over the class of C^∞ coefficients, even for "large" times. Additionally, we will show that the wave atom matrix has a natural *low-rank block structure* after separation of the spatial indices.

In Section 4.1 we give a more precise description of wave atoms, including a fast digital transform. Section 4.2 introduces the repeated squaring algorithm, its refinement involving the separated wave atom representation, and complexity estimates. We present numerical experiments in Section 4.3. The rest of the chapter is then devoted to stating and proving estimates of ϵ-separation ranks in various situations.

4.1 Wave Atoms

In this section we show how to implement wave atoms as a fast digital transform. The requirements we put on a family of basis function to be called "wave atoms" are quite stringent and will be made precise in Section 4.1.1. They have to do with uniform space-frequency localization, and put the general architecture introduced in Section 1.2.4 on solid ground. To the best of our knowledge, none of the transforms in the repertoire of modern computational harmonic analysis satisfies these localization conditions.

4.1.1 Definition of Wave Atoms

We write wave atoms as $\varphi_\mu(x)$, with subscript $\mu = (j, \mathbf{m}, \mathbf{n}) = (j, m_1, m_2, n_1, n_2)$. All five quantities j, m_1, m_2, n_1, n_2 are integer-valued and index a point (x_μ, ξ_μ) in phase-space, as

$$x_\mu = 2^{-j}\mathbf{n}, \qquad \xi_\mu = \pi 2^j \mathbf{m}, \qquad C_1 2^j \le \max_{i=1,2} |m_i| \le C_2 2^j. \tag{4.1}$$

where C_1, C_2 are two positive constants left unspecified for convenience, but whose values will be implied by the specifics of the implementation. Heuristically, the *position vector* x_μ is the center of $\varphi_\mu(x)$ and the *wave vector* ξ_μ determines the centers of both bumps of $\hat{\varphi}_\mu(\xi)$ as $\pm\xi_\mu$. Note that the range of $\mathbf{m}$ needs to be further reduced to $m_2 > 0$, (or $m_2 = 0$ and $m_1 > 0$,) to account for the central symmetry of the Fourier transform of real-valued functions about the origin in ξ. Some further restriction on $\mathbf{n}$ (cutoff in space) and j (cutoff in scale), are of course necessary in practice, but not for the description of a frame of L^2.

Wave atoms then need to obey a localization condition around the phase-space point (x_μ, ξ_μ).

Definition 4.1. *(Wave Atoms) Let x_μ and ξ_μ be as in equations (4.1) for some C_1, C_2. The elements of a frame of wave packets $\{\varphi_\mu\}$ are called* wave atoms *when*

$$|\hat{\varphi}_\mu(\xi)| \le C_M \cdot 2^{-j}(1+2^{-j}|\xi-\xi_\mu|)^{-M} + C_M \cdot 2^{-j}(1+2^{-j}|\xi+\xi_\mu|)^{-M}, \qquad \textit{for all } M > 0, \tag{4.2}$$

and

$$|\varphi_\mu(x)| \le C_M \cdot 2^j (1 + 2^j |x - x_\mu|)^{-M}, \qquad \textit{for all } M > 0. \tag{4.3}$$

It is of course possible to restrict the decay order or even moderately alter the definition of x_μ and ξ_μ—and still call the basis functions "wave atoms"—but this is a refinement we will not address here.

The parabolic scaling is encoded in the localization conditions as follows: at scale 2^{-2j}, or frequency $|\xi_\mu| \sim 2^{2j}$, the essential frequency support is of size $\sim 2^j$ (for each bump) and the essential spatial support is of size $\sim 2^{-j}$. Note that the subscript j indexes the different "dyadic coronae," whereas the additional subscript $\mathbf{m}$ labels the different wave numbers ξ_μ within each dyadic corona.

4.1.2 Properties

In the same spirit as curvelets, there is a natural notion of pseudodistance in phase-space associated to wave atoms.

Definition 4.2. *Let $\mu = (j, \mathbf{m}, \mathbf{n})$ and $\mu' = (j', \mathbf{m}', \mathbf{n}')$ be two wave atom subscripts. The* wave atom pseudodistance $\omega(\mu, \mu')$ *is defined as*

$$\omega(\mu, \mu') = 1 + 2^{\min(j,j')}|x_\mu - x_{\mu'}| + 2^{-\max(j,j')}|\xi_\mu - \xi_{\mu'}|.$$

The motivation for this definition is the interpretation of ω a *lattice distance* in phase-space. Consider the graph with wave packet indices μ as nodes, and connection between two nodes if and only if the corresponding wave packets are neighbors, that is if

- either $j = j'$, $\mathbf{m} = \mathbf{m}'$ and $|\mathbf{n} - \mathbf{n}'| = 1$;
- or μ and μ' correspond to adjacent frequency tiles and $|x_\mu - x_{\mu'}| = \min_{\mu''=(j',\mathbf{m}',\mathbf{n}'')} |x_\mu - x_{\mu''}|$.

Then $\omega(\mu, \mu')$ is proportional to the minimum number of edges needed to connect μ and μ'.

Of course ω is not a distance in the strict sense, but it is symmetric, satisfies the quasi-triangle inequality $\omega(\mu, \mu'') \le C(\omega(\mu, \mu') + \omega(\mu', \mu''))$, and is invariant under Hamiltonian flows, $\omega(\mu(t), \mu'(t)) \asymp \omega(\mu(0), \mu'(0))$.

The main purpose of ω is that it allows us to formulate a key almost-orthogonality estimate.

Lemma 4.1. *Let φ_μ and $\tilde{\varphi}_\mu$ be two collections of wave atoms, in the sense of Definition 4.1. Then for every $M > 0$ there exists a constant $C_M > 0$ such that*

$$|\langle \varphi_\mu, \tilde{\varphi}_{\mu'} \rangle| \le C_M \cdot \omega(\mu, \mu')^{-M}.$$

Proof. We will make use of the following elementary bump convolution convolution inequality (see [65] p. 56), valid when $a \ge a'$ and, say, $M \ge 2$,

$$\int_{-\infty}^{\infty} (1 + a|x|)^{-M}(1 + a'|x - x_0|)^{-M}\, dx \le \frac{C}{a}(1 + a'|x_0|)^{-M}. \tag{4.4}$$

Assume without loss of generality that $j \leq j'$. Combining (4.4) with the frequency localization estimate (4.2), we obtain

$$\int |\hat{\varphi}_\mu(\xi)\hat{\tilde{\varphi}}_{\mu'}(\xi)|\, d\xi \leq C_M \cdot 2^{-j'}(1 + 2^{-j'}|\xi_\mu - \xi_{\mu'}|)^{-M}. \tag{4.5}$$

Similarly, combining (4.4) with the spatial localization estimate (4.3), we obtain

$$\int |\varphi_\mu(x)\tilde{\varphi}_{\mu'}(x)|\, dx \leq C_M \cdot 2^{j}(1 + 2^{j}|x_\mu - x_{\mu'}|)^{-M}. \tag{4.6}$$

The conclusion follows by taking the geometric mean of (4.5) and (4.6), and noticing that

$$(1 + 2^{-j'}|\xi_\mu - \xi_{\mu'}|)^{-M}(1 + 2^{j}|x_\mu - x_{\mu'}|)^{-M} \leq (1 + 2^{-j'}|\xi_\mu - \xi_{\mu'}| + 2^{j}|x_\mu - x_{\mu'}|)^{-M}.$$

□

Almost orthogonality is one ingredient in the proof of the main sparsity result, theorem 1.1, hence the resemblance of statements.

4.1.3 Implementation of Wave Atoms: 1D Warmup

In practice, wave atoms will be constructed from tensor products of adequately chosen 1D wave packets.

We will first build a one-dimensional family of real-valued wave packets $\psi^j_{m,n}(x)$, $j \geq 0, m \geq 0, n \in \mathbb{Z}$, centered in frequency around $\pm\,\xi_{j,m} = \pm\,\pi 2^j m$, with $C_1 2^j \leq m \leq C_2 2^j$; and centered in space around $x_{j,n} = 2^{-j}n$. The one-dimensional version of the parabolic scaling dictates that the support of (each bump of) $\hat{\psi}^j_{m,n}(\xi)$ be of length $O(2^j)$ while $\xi_{j,m} = O(2^{2j})$. The desired corresponding tiling of frequency is illustrated at the bottom of Figure 4.1.

Filterbank-based wavelet packets naturally come to mind as a potential definition of an orthonormal basis satisfying these localization properties. They also come with fast algorithms. The wavelet packet tree, defining the partitioning of the frequency axis in 1D, can be chosen to have depth j when the frequency is $\sim 2^{2j}$, as illustrated in Figure 4.1. However, there is a well-documented problem associated with standard wavelet packets, namely that the sense in which they satisfy frequency localization is rather weak. It is an unavoidable feature of the filterbank architecture that the uncertainty (product of time

and frequency deviations) increases with the frequency, instead of remaining close to the Heisenberg bound. For references for precise estimates of the "wavelet packet curse," see [69, 88]. As a result, in our context, we cannot hope to satisfy the wave atom definition using basis functions which come from a wavelet packet analysis.

An elegant solution to the frequency localization problem has been given by Lars Villemoes in [88]. The trick consists in exhibiting adequate symmetric pairs of compactly supported bumps in frequency, given by the formula

$$\hat{\psi}^0_m(\xi) = e^{-i\xi/2}\left[e^{i\alpha_m} g(\epsilon_m(\xi - \pi(m+\frac{1}{2}))) + e^{-i\alpha_m} g(\epsilon_{m+1}(\xi + \pi(m+\frac{1}{2})))\right],$$

where $\epsilon_m = (-1)^m$ and $\alpha_m = \frac{\pi}{2}(m+\frac{1}{2})$. The function g is an appropriate real-valued, C^∞ bump function, compactly supported on an interval of length 2π, and chosen such that

$$\sum_m |\hat{\psi}^0_m(\xi)|^2 = 1.$$

Then the translates $\psi_m(t-n)$ form an orthonormal basis of $L^2(\mathbb{R})$. This construction provides a *uniform* tiling of the frequency axis,[1] in the sense that every bump in frequency has the same support size, 2π.

Multiscale tilings like the one in Figure 4.1 can be obtained by combining dyadic dilates and translates of $\hat{\psi}^0_m$ on the frequency axis. We need to introduce the subscript j to index scale, and write our basis functions as $\psi^j_{m,n}(x)$. To preserve orthonormality of the $\psi^j_{m,n}(x)$, the profile g needs to be asymmetric in addition to all the other properties, in the sense that

$$g(-2\xi - \frac{\pi}{2}) = g(\frac{\pi}{2} + \xi)$$

for $\xi \in [-\pi/3, \pi/3]$, with g itself supported on $[-7\pi/6, 5\pi/6]$. We say $g(\xi)$ is "left-handed", whereas $g(-\xi)$ is "right-handed". As a result, the uniform partitioning of the frequency axis is obtained as an alternating sequence of staggered left-handed and right-handed bumps. A scale doubling can be achieved by concatenating two left-handed bumps at scales differing by a factor 2. Figure 4.1, bottom row, depicts the Villemoes system where three scale doublings have been implemented and marked by crosses. Remark in passing that scale halving could be implemented using right-handed bumps at different scales, but scale quadrupling is

[1] For specialists, the secret to avoiding the Balian-Low no-go theorem is to use *two* bumps in frequency.

impossible using the present scheme.

As for labeling, note that the couple (j, m) refers to a point on the wavelet packet tree; the depth at that point is $J - j$, where J is the maximum depth, and m can be interpreted as the number of nodes on the left at the same depth (nodes are not necessarily leaves). This is the standard indexing scheme for wavelet packets used in [62]. The translation step is now 2^{-j} at scale j, whereas each bump in frequency is supported on an interval of length $2^j \times 2\pi$. Choosing the "parabolic tree" as in Figure 4.1 amounts to specifying the wave vector $\xi_{j,m}$, defined as the center of the positive frequency bump, as

$$\xi_{j,m} = \pi 2^j m \sim 2^{2j}$$

The resulting basis of wavelet packets $\{\psi_m^j(x - 2^{-j}n)\}$ is orthonormal for $L^2(\mathbb{R})$.

The implementation of the Villemoes wavelet packets is rather straightforward in the frequency domain. For each wave number $\xi_{j,m}$, the coefficients $c_{j,m,n}$ can be seen as a decimated convolution at scale 2^{-j},

$$c_{j,m,n} = \int \psi_m^j(x - 2^{-j}n)u(x)\,dx,$$

By Plancherel,

$$c_{j,m,n} = \frac{1}{2\pi}\int e^{-i2^{-j}n\xi}\hat{\psi}_m^j(\xi)\hat{u}(\xi)\,d\xi$$

Assuming that the function u is accurately discretized at $x_k = kh$, $h = 1/N$, $k = 1, \ldots, N$, then up to some small truncation error,

$$c_{j,m,n} \simeq \frac{1}{2\pi}\sum_{k=2\pi(-N/2+1:1:N/2)} e^{-i2^{-j}nk}\hat{\psi}_m^j(k)\hat{u}(k). \tag{4.7}$$

If the data $\hat{\psi}_m^j(k)\hat{u}(k)$ were supported in an interval of length $2^j \times 2\pi$, then the above sum could be restricted to values of k inside that interval, and computed efficiently using a reduced inverse FFT, of size 2^j. In reality the support properties of g are such that the data is supported inside *two* disjoints intervals of size $2^{j+1}\pi$, symmetric about the origin. The sum (4.7) can still be computed by a reduced inverse FFT provided $\hat{\psi}_m^j\hat{u}$ is folded by $2^{j+1}\pi$-periodicity inside an interval of size $2^{j+1}\pi$ centered about the origin. This trick was already used for the implementation of the discrete curvelet transform and is called

wrapping, see chapter 3. The simple algorithm for wavelet packets is then the following.

- Perform a FFT of size N of the samples $u(x_k)$.
- For each pair (j, m), wrap the product $\hat{\psi}^j_m \hat{u}$ by periodicity inside the interval $[-2^j\pi, 2^j\pi]$. Then perform an inverse FFT of size 2^j of the result to obtain $c_{j,m,n}$.
- Repeat over (j, m).

The complexity of each inverse FFT at scale j is $O(j2^j)$, and there are $O(2^j)$ frequency bumps at scale j, indexed by m, so the total complexity is

$$\sum_{0 \leq j \leq J} O(j2^{2j}) = O(J2^{2J}) = O(N \log N),$$

with $N = 2^{2J}$.

Since the $c_{j,m,n}$ are coefficients in an orthonormal basis, the inverse transform is simply the adjoint and can be computed by reversing all the steps in the above algorithm.

- For each (j, m), perform a FFT in n of $c_{j,m,n}$, of size 2^j, then unwrap the result on the frequency axis around the support of $\hat{\psi}^j_m$.
- Sum the contributions corresponding to all the couples (j, m).
- Perform an inverse FFT, of size N, to obtain $u(x_k)$.

Likewise, the complexity of the inverse transform is $O(N \log N)$.

The decomposition into two bumps, of positive and negative frequency respectively, can be written

$$\hat{\psi}^j_{m,n}(\xi) = \hat{\psi}^j_{m,n,+}(\xi) + \hat{\psi}^j_{m,n,-}(\xi) \tag{4.8}$$

with the symmetry relation $\hat{\psi}^j_{m,n,-}(\xi) = \overline{\hat{\psi}^j_{m,n,+}(-\xi)}$ which owes to the real-valuedness of $\psi^j_{m,n}$. After preforming a Hilbert transform, $H\psi^j_{m,n}$ is another orthonormal basis of $L^2(\mathbb{R})$. In the frequency domain, Hilbert transformation amounts to taking a linear combination of the two bumps with weights $(-i, i)$ instead of $(1, 1)$:

$$\widehat{H\psi^j_{m,n}}(\xi) = -i\hat{\psi}^j_{m,n,+}(\xi) + i\hat{\psi}^j_{m,n,-}(\xi). \tag{4.9}$$

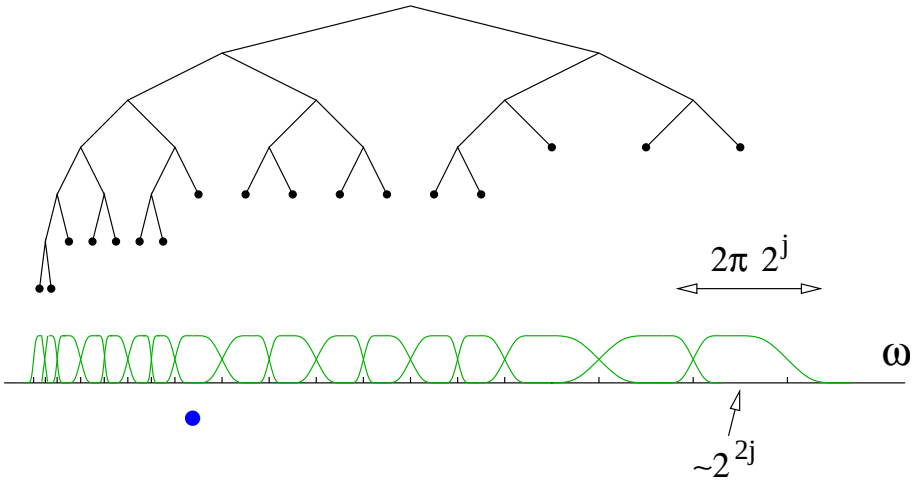

Figure 4.1: The wavelet-packet tree corresponding to wave atoms. Textbook material on wavelet packet trees can be found in [62], chapter 8. The bottom graph depicts Villemoes wavelet packets on one half of the frequency axis. The dot indicates the frequency where a change of scale occurs.

Even though we made a point that the Villemoes basis functions do not come from a standard multiresolution analysis, we still continue to call them "wavelet packets" in what follows.

Finally, we would like to mention Eric Kolaczyk's approach [55] for the implementation of Meyer wavelets, and by extension Villemoes wavelet packets, based on local cosine windowing of the frequency axis. His approach explains in a transparent manner the aliasing cancellations during inversion of the wavelet transform, though the implementation looks involved in contrast to the "wrapping" approach.

4.1.4 Implementation of Wave Atoms: 2D Extension

Two-dimensional orthonormal basis functions with 4 bumps in the frequency plane can be formed by individually taking products of 1D wavelet packets. As we did in earlier sections, let us abbreviate $\mu = (j, \mathbf{m}, \mathbf{n})$, where $\mathbf{m} = (m_1, m_2)$ and $\mathbf{n} = (n_1, n_2)$. We write

$$\varphi_\mu^+(x_1, x_2) = \psi_{m_1}^j(x_1 - 2^{-j} n_1)\, \psi_{m_2}^j(x_2 - 2^{-j} n_2),$$

The Fourier transform is also separable, namely

$$\hat{\varphi}_\mu^+(\xi_1, \xi_2) = \hat{\psi}_{m_1}^j(\xi_1) e^{-i 2^j (n_1 x_1)}\, \hat{\psi}_{m_2}^j(\xi_2)\, e^{-i 2^j (n_2 x_2)}.$$

A dual orthonormal basis can be defined from the Hilbert-transformed wavelet packets,

$$\varphi_\mu^-(x_1, x_2) \equiv H\psi_{m_1}^j(x_1 - 2^{-j}n_1)\ H\psi_{m_2}^j(x_2 - 2^{-j}n_2).$$

It is easy to see from relations (4.8) and (4.9) that the recombination

$$\varphi_\mu^{(1)} = \frac{\varphi_\mu^+ + \varphi_\mu^-}{2}, \qquad \varphi_\mu^{(2)} = \frac{\varphi_\mu^+ - \varphi_\mu^-}{2},$$

provides basis functions with *two* bumps in the frequency plane, symmetric with respect to the origin, hence directional wave packets. Together, $\varphi_\mu^{(1)}$ and $\varphi_\mu^{(2)}$ form the wave atom frame and may be denoted jointly as φ_μ. The price to pay in considering both $\varphi_\mu^{(1)}$ and $\varphi_\mu^{(2)}$ is an increase of a factor 2 in the redundancy. Wave atoms otherwise remain a tight frame, in the sense that

$$\sum_\mu |\langle \varphi_\mu^{(1)}, u\rangle|^2 + \sum_\mu |\langle \varphi_\mu^{(2)}, u\rangle|^2 = \|u\|^2.$$

They also satisfy all the wave atom properties, by construction.

As the reader might have noticed, the construction is not a simple tensor product because there is only one scale subscript j. This is akin to the construction of "nonstandard," or MRA wavelet bases in 2D where the point is to enforce the same scale in both directions, hence an isotropic aspect ratio. The resulting tiling of the frequency plane is shown in Figure 4.2.

In practice, the algorithm for wave atoms is based on the obvious generalization of the 1D wrapping strategy to two dimensions – except for a slight complication. The admissible tilings of the frequency plane at scale j are restricted by

$$\max_{i=1,2} |m_i| = 4n_j + 2,$$

for some integer n_j depending on j. In Figure 4.2, we check that this property holds with $n_0 = 0$, $n_1 = 1$ and $n_2 = 2$. The rationale for this (benign) restriction is that a window needs to be left-handed in both directions near a scale doubling (for positive frequencies), and that this parity needs to match with the rest of the lattice. The rule is that $\hat{\psi}_{m,+}^j$ is left-handed for m even and right-handed for m odd, so for instance $\hat{\psi}_3^2(\xi_1)\hat{\psi}_3^2(\xi_2)$ would not be an admissible window near a scale doubling, whereas $\psi_4^2(x_1)\psi_4^2(x_2)$ is admissible (and

its tile is indicated in Figure 4.2 by a dot).

In complete analogy with the 1D case, the complexity of wave atoms is of course $O(N^2 \log N)$.

The adjoint wave atom transform is still an inverse (actually the Moore-Penrose pseudo-inverse) because the frame is tight. As explained previously, it is computed by reversing all the operations in the direct transform, and also takes $O(N^2 \log N)$ operations.

Note that the same 2D recombination strategy, but based on standard wavelet packets instead of Villemoes packets, has already been considered in the field of image processing under the name "dual-tree M-band wavelets," see [24].

Let us conclude this section by mentioning that, in the present implementation, most wave atoms have an infinite number of (directional) *vanishing moments* in the sense that

$$\int \varphi_\mu(y_1, y_2) y_1^k \, dy_1 = 0, \qquad \text{for all } \; k \geq 0, \qquad \text{for all } \; y_2, \tag{4.10}$$

where y_1 is the spatial coordinate along the wave vector ξ_μ, and y_2 is perpendicular. This property follows from the fact that most atoms $\hat{\varphi}_\mu(\xi)$ vanish in a (large) strip including the origin, oriented in the direction perpendicular to ξ_μ. In contrast, those few atoms obeying $\hat{\varphi}_\mu(0) \neq 0$ form the small minority of "coarse scale" wave atoms, when $j = 0$.

4.1.5 The Orthobasis Variation

In practice, one may want to work with an orthonormal basis instead of a tight frame to represent the wave equation, for example $\varphi_\mu^+(x)$. The consequence of this choice is that each basis functions would oscillate in *two* distinct directions, instead of one.

4.2 Main Algorithm

The description of the main algorithm will be split into two parts. The basic repeated squaring algorithm only exploits sparsity of the wave atom matrix of the propagator $E(t)$ and is detailed in Section 4.2.1. The refinement based on separation of spatial indices comes as a modification of the basic repeated squaring algorithm, and is explained in Section 4.2.2.

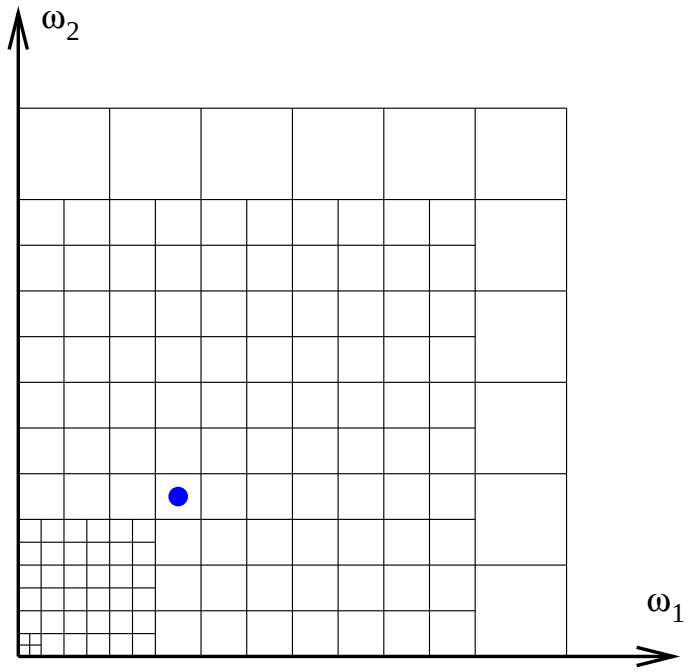

Figure 4.2: The wave atom tiling of the frequency plane. Only the first quadrant is shown. The dot below the ω axis indicates the same change of scale as in Figure 4.1 and corresponds to the basis function denoted $\psi_4^2(x_1)\psi_4^2(x_2)$ in the text.

4.2.1 Basic Repeated Squaring

Let us denote $u(t)$ for the couple $(p(t), \frac{\partial p}{\partial t})$, and write the wave equation as the first-order system $\frac{\partial u}{\partial t} = Au$ with initial condition $u(0) = u_0$. The generator is

$$A = \begin{pmatrix} 0 & I \\ c^2(x)\Delta & 0 \end{pmatrix}. \tag{4.11}$$

We define the propagator $E(t)$ from $u(t) = E(t)u_0 = e^{tA}u_0$.

Since the solution $u(t)$ has two components, we need to introduce $\mathbf{e}_1 = (1, 0)$ and $\mathbf{e}_2 = (0, 1)$. The wave atom matrix elements are

$$E(t; \mu\nu; \mu'\nu') = \langle E(t)\varphi_{\mu'}\mathbf{e}_{\nu'}, \varphi_\mu \mathbf{e}_\nu \rangle.$$

We write $\tilde{E}(t; \mu\nu; \mu'\nu')$ for its numerical approximation. As mentioned in the introduction we aim at building this matrix at dyadic times $t_n = 2^n \Delta t$, using a repeated-squaring

strategy, based on the group property

$$E(2t;\mu\nu;\mu'\nu') = \sum_{\mu'',\nu''} E(t;\mu\nu;\mu''\nu'')E(t;\mu''\nu'';\mu'\nu'),$$

which in turn comes from $E(2t) = (E(t))^2$ and the tight-frame property. The squaring is efficient because the numerical approximation of the wave atom matrices is kept sparse at all dyadic times, by putting to zero the small entries below a prescribed threshold.

Algorithm 4.1. *(Wave-Atom Repeated Squaring) Choose a small time step Δt and a small tolerance ϵ. Denote by* Trunc *the operation of putting to zero all matrix elements below ϵ in absolute value.*

- **Initialization***: Obtain $\tilde{A}(\mu\nu;\mu'\nu')$ an approximation to the wave atom matrix of the generator A, then*

$$\tilde{E}(\Delta t,\mu\nu;\mu'\nu') = \delta_{\mu\nu;\mu'\nu'} + \Delta t\ Trunc(\tilde{A}(\mu\nu;\mu'\nu')).$$

- **Iteration***: Forecast the biggest entries' location, then compute them as*

$$\tilde{E}(2^{n+1}\Delta t;\mu\nu;\mu'\nu') = Trunc \sum_{\mu'',\nu''} E(2^n\Delta t;\mu\nu;\mu''\nu'')E(2^n\Delta t;\mu''\nu'';\mu'\nu').$$

- **Terminate** *at time $\tau = 2^{n^*}\Delta t$.*

To compute the solution $u(\tau)$ at time τ, start with the coefficients

$$c_{\mu\nu}(0) = \langle u_0, \varphi_\mu \mathbf{e}_\nu\rangle,$$

perform the matrix-vector multiplication,

$$\tilde{c}_{\mu\nu}(\tau) = \sum_{\mu'\nu'} \tilde{E}(\tau;\mu\nu;\mu'\nu')c_{\mu'\nu'}(0),$$

and inverse transform, $\tilde{u}(\tau) = \sum_{\mu\nu} \tilde{c}_{\mu\nu}(\tau)\varphi_\mu \mathbf{e}_\nu$. For times larger than τ, one should perform several applications of $E(\tau)$ to the initial condition.

As alluded to in the introduction section, prediction of the location of the large matrix elements is efficiently done using the Phase-Flow Method (PFM), see [90]. The truncation

should be done to keep at most B elements per row (hence also per column), with B a moderately large constant.

For the initialization, how to compute an approximation to the generator $A_{\mu\nu;\mu'\nu'}$ in an efficient manner is best understood in the context of the separated wave atom representation, so we postpone this discussion until the next section.

Let us finally remark that, in view of Theorem 1.1, the *Trunc* operation consists in keeping track of *two* shifted diagonals, because there are two Hamiltonian flows. If instead of the standard wave atoms $\varphi_\mu^{(1)}$ or $\varphi_\mu^{(2)}$ we use the orthobasis variation of Section 4.1.5, $\varphi_m u^+ = \varphi_\mu^{(1)} + \varphi_\mu^{(2)}$, we could expect to have to trace *four* bumps. For small times, namely $t \leq \frac{1}{\sqrt{N}}$ or a multiple thereof, we will see later in Section 4.2.3 that tracing is useful but not necessary, so the gain due to lower redundancy may offset the frequency entangling. For larger times, "clean" wave atoms with two bumps in frequency may be more appropriate.

4.2.2 The Separated Wave Atom Representation

In the wave atom representation of $E(t)$, consider the submatrix left after fixing ν, ν' and the wave vectors $(j, \mathbf{m})$ and $(j', \mathbf{m}')$. The remaining indices are those of the position vectors $\mathbf{n} = (n_1, n_2)$ and $\mathbf{n}' = (n_1', n_2')$. The *separated wave atom representation* is obtained by seeking a low-rank approximation corresponding to separation of the spatial indices along x_1 vs. x_2,

$$E(t; j, \mathbf{m}, \mathbf{n}, \nu; j', \mathbf{m}', \mathbf{n}', \nu') = \sum_{k=1}^{r} \sigma_k u^k_{n_1,n_1'} v^k_{n_2,n_2'} + O(\epsilon), \tag{4.12}$$

where u^k and v^k have been normalized to unit ℓ_2 norm. Of course u^k and v^k depend on $j, \mathbf{m}, \nu; j', \mathbf{m}', \nu'$. The most efficient such decomposition, in the sense that the ℓ^2 norm of the residual is minimized for fixed r, is the singular value decomposition of the block $(j, \mathbf{m}, \nu; j', \mathbf{m}', \nu')$ after *reorganization* of the matrix elements to make the row and column indices $(n_1, n_1'; n_2, n_2')$ instead of $(n_1, n_2; n_1', n_2')$.

Conversion from the standard to the separated wave atom representation, as an SVD factorization of the reorganized submatrix, is however never done in practice. Instead, we modify the repeated squaring strategy so that all computations are done on separated components without ever forming the standard submatrix. Let us explain how both initialization and matrix multiplication can be realized in this context.

4.2.2.1 Initialization

The wave atom representation of the generator reads $\langle \varphi_{\mu\nu}, A\varphi_{\mu'\nu'} \rangle$, where $\nu = 1$ refers to the $p(t)$ component, whereas $\nu = 2$ refers to the $\frac{\partial p}{\partial t}$ component. The only non-trivial or non-precomputable contribution is

$$\langle \varphi_{\mu,2}, A\varphi_{\mu',1} \rangle = \int_{[0,1]^2} \varphi_\mu(x) c^2(x) \Delta \varphi_{\mu'}(x)\, dx. \tag{4.13}$$

Our initialization strategy is based on separation of the integrand in x_1 vs. x_2. Since $c^2(x)$ is a C^∞ periodic function, its ϵ-separation rank is a small constant $C_\epsilon = O(\epsilon^{-1/M})$ for all $M > 0$ (see Lemma 4.2) so we can write

$$c^2(x) = \sum_{k=1}^{C_\epsilon} \gamma_k^{(1)}(x_1) \gamma_k^{(2)}(x_2).$$

The Laplacian operator is also nicely separated,

$$\Delta = \frac{\partial^2}{\partial x_1^2} \otimes I + I \otimes \frac{\partial^2}{\partial x_2^2}. \tag{4.14}$$

As for wave atoms themselves, let us assume that we are using the separable "orthobasis" variation, as in Section 4.1.5. For the full wave atoms there would be two separated terms to write down instead.

We can then split the matrix element (4.13) into a finite number of separated components,

$$\begin{aligned} \langle \varphi_\mu, c^2(x) \Delta \varphi_{\mu'} \rangle = \sum_{k=1}^{C_\epsilon} & \langle \psi^j_{m_1,n_1}, \gamma_k^{(1)} \frac{\partial^2}{\partial x_1^2} \psi^{j'}_{m_1',n_1'} \rangle \, \langle \psi^j_{m_2,n_2}, \gamma_k^{(2)} \psi^{j'}_{m_2',n_2'} \rangle \\ & + \langle \psi^j_{m_1,n_1}, \gamma_k^{(1)} \psi^{j'}_{m_1',n_1'} \rangle \, \langle \psi^j_{m_2,n_2}, \gamma_k^{(2)} \frac{\partial^2}{\partial x_2^2} \psi^{j'}_{m_2',n_2'} \rangle, \end{aligned}$$

where all the inner products in the right-hand side are one-dimensional. Observe that the above formula is exactly in "separated wave atom" form, as in equation (4.12).

The initialization algorithm computes all the factors in the above decomposition as follows. Assume the segment $[0, 1]$ has been discretized into N equispaced points. For each (j', m_1', n_1'),

1. Form $\psi^{j'}_{m_1',n_1'}(x_1)$ for x_1 on the grid by applying the inverse 1D wavelet packet trans-

form to the sequence of coefficients

$$c_{j,m_1,n_1} = \begin{cases} 1 & \text{if } j = j', m_1 = m_1', n_1 = n_1', \\ 0 & \text{otherwise.} \end{cases}$$

2. Apply $\gamma_k^{(1)} \frac{\partial^2}{\partial x_1^2}$ to $\psi^{j'}_{m_1',n_1'}(x)$. For accuracy purposes, all derivatives are discretized in the Fourier domain.

3. Apply a direct wavelet packet transform to the result and obtain at once the inner products with all the $\psi^j_{m_1,n_1}(x_1)$.

Repeat over all indices (j', m_1', n_1'). Repeat the algorithm, *mutatis mutandis*, for the inner products involving no derivatives and the inner products involving x_2 instead of x_1. Do not sum over k or multiply the factors as we are interested in the separated form only.

In practice, we may use a different time-integration scheme than Euler explicit for the first time step. We have found the leap-frog scheme to be quite efficient.

4.2.2.2 Matrix Multiplication

We seek a fast algorithm for

$$E(2t; j, \mathbf{m}, \mathbf{n}, \nu; j'', \mathbf{m}'', \mathbf{n}'', \nu'') = \sum_{j', \mathbf{m}', \mathbf{n}', \nu'} E(t; j, \mathbf{m}, \mathbf{n}, \nu; j', \mathbf{m}', \mathbf{n}', \nu') E(t; j', \mathbf{m}', \mathbf{n}', \nu'; j'', \mathbf{m}'', \mathbf{n}'', \nu''), \tag{4.15}$$

where each factor is given by (4.12). We fix the row index $(j, \mathbf{m}, \nu)$ as well as the column index $(j'', \mathbf{m}'', \nu'')$, and for simplicity omit to write them in what follows.

We start by directly computing each sum over n_1' and n_2'. There is one such sum for each value of the intermediate index $(j', \mathbf{m}', \nu')$. Let us introduce

$$U^{\mathbf{k}}_{n_1,n_1''}(j', \mathbf{m}', \nu') = \sum_{n_1'} u^k_{n_1,n_1'} u^{k'}_{n_1',n_1''} \tag{4.16}$$

and, similarly,

$$V^{\mathbf{k}}_{n_2,n_2''}(j', \mathbf{m}', \nu') = \sum_{n_2'} u^k_{n_2,n_2'} u^{k'}_{n_2',n_2''}. \tag{4.17}$$

We grouped (k, k') into one single index $\mathbf{k}$. If we also let $\sigma_{\mathbf{k}} = \sigma_k \sigma_{k'}$, and form the resulting

diagonal matrix $\Sigma_{\mathbf{k}}$, then the matrix element (4.15) can be written as

$$\sum_{j',\mathbf{m}',\nu'} U(j',\mathbf{m}',\nu')\,\Sigma\,V^t(j',\mathbf{m}',\nu') \equiv \sum_{\mathbf{k}} \sum_{j',\mathbf{m}',\nu'} U^{\mathbf{k}}_{n_1,n_1''}(j',\mathbf{m}',\nu')\sigma_{\mathbf{k}} V^{\mathbf{k}}_{n_2,n_2''}(j',\mathbf{m}',\nu'), \quad (4.18)$$

Call K the maximum number of different values of the couple $\mathbf{k}$; if $k \le r$, then $K \le r^2$. Call M the maximum number of different values of the triple $(j',\mathbf{m}',\nu')$. To obtain the desired separated wave atom form we need not only compute those sums, but also factor the result into its singular value decomposition,

$$\tilde{U}\,\tilde{\Sigma}\,\tilde{V}^t \equiv \sum_{\tilde{k}} \tilde{U}^{\tilde{k}}_{n_1,n_1''}\tilde{\sigma}_{\tilde{k}}\tilde{V}^{\tilde{k}}_{n_2,n_2''}, \quad (4.19)$$

where $\tilde{U}$ and $\tilde{V}$ are isometric matrices.

So we are faced with the problem of computing the SVD of a sum of matrices which are almost in SVD form—because the columns of $U(j',\mathbf{m}',\nu')$ and $V(j',\mathbf{m}',\nu')$ are not in general orthogonal. We will proceed in two steps:

- We start by turning each U and V into isometric matrices. For each $(j',\mathbf{m}',\nu')$, perform a QR decomposition to obtain an isometric matrix $Q_U(j',\mathbf{m}',\nu')$ and an upper triangular matrix $R_U(j',\mathbf{m}',\nu')$ such that

  $$U(j',\mathbf{m}',\nu') = Q_U(j',\mathbf{m}',\nu')R_U(j',\mathbf{m}',\nu').$$

  Similarly, factor

  $$V(j',\mathbf{m}',\nu') = Q_V(j',\mathbf{m}',\nu')R_V(j',\mathbf{m}',\nu').$$

  Gather the small factors in the middle and perform an SVD:

  $$R_U(j',\mathbf{m}',\nu')\;\Sigma\;R_V^t(j',\mathbf{m}',\nu') = U^\sharp(j',\mathbf{m}',\nu')\;\Sigma^\sharp(j',\mathbf{m}',\nu')\;[V^\sharp]^t(j',\mathbf{m}',\nu').$$

  Put to zero the small singular values in $\Sigma^\sharp(j',\mathbf{m}',\nu')$ below some threshold τ, so as to keep at most $O(r)$ of them. Then compute

  $$U(j',\mathbf{m}',\nu') := Q_U(j',\mathbf{m}',\nu')\,U^\sharp(j',\mathbf{m}',\nu'),$$

$$V(j', \mathbf{m}', \nu') := Q_V(j', \mathbf{m}', \nu')\, V^\sharp(j', \mathbf{m}', \nu'),$$

and for simplicity call $\Sigma(j', \mathbf{m}', \nu') := \Sigma^\sharp(j', \mathbf{m}', \nu')$. We have just orthogonalized (4.18) at the (benign) expense of making each Σ matrix depend on $(j', \mathbf{m}', \nu')$.

- We can now simplify the sum over $(j', \mathbf{m}', \nu')$. Let us group terms two by two and notice that a sum of two SVDs can be rewritten in matrix form as

$$U_1\, \Sigma_1\, V_1^t + U_2\, \Sigma_2\, V_2^t = \begin{pmatrix} U_1 & V_1 \end{pmatrix} \begin{pmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{pmatrix} \begin{pmatrix} V_1^t \\ V_2^t \end{pmatrix}. \tag{4.20}$$

  The same strategy as above, involving two QR and one SVD decomposition, can be invoked to compute the SVD of the right-hand side in the above equation. The procedure can be applied to each couple of terms and repeated at the next level. This way, the whole sum (4.18) can be reduced in a binary fashion into its SVD form, leaving us with (4.19).

Standard linear algebra routines have been used for QR and SVD. It does not appear that iterative algorithms for sparse SVD offer any improvement, in the present context, over the standard algorithms.

4.2.2.3 Upscaled Timestepping

Once the separated wave atom representation of the wave propagator at time τ is available, we can apply it to the initial condition as follows.

- Apply the wave atom transform to each component of the initial condition u_0.
- For each $j, \mathbf{m}, \nu$ and $j', \mathbf{m}', \nu'$, unfold the separated form to obtain the classical wave atom representation. Not all matrix elements in the classical form need to be computed, however. For a given (n_1, n_2), only a certain subset of positions subscripts (n_1', n_2') will be relevant – only those for which the wave has been given enough time to travel from x_μ to $x_{\mu'}$.
  - For times $\tau \lesssim 2^{-j}$, a wave atom cannot travel essentially farther than its own diameter, hence the restriction $|\mathbf{n}' - \mathbf{n}| \leq C$ for some constant C to be determined empirically.

  - For times $\tau \gtrsim 2^{-j}$, this rule becomes $|\mathbf{n}' - \mathbf{n}| \leq C \cdot t 2^j$.

  Once the restricted submatrices have been formed, we can compute the large matrix-vector product of the propagator with the coefficients of the initial condition.

- Apply an inverse wave atom transform to get $u(t)$ from its coefficients.

The same procedure, without the initial and final transforms, can be iterated to perform timestepping with the "upscaled" time step τ, generally much larger than the CFL timestep. As we will see from the complexity analysis in the next section, a reasonable choice of time step is $\tau \simeq \frac{1}{\sqrt{N}}$, whereas the CFL timestep is at most $\Delta t < \frac{1}{c_{\max} N}$ when $c(x) \leq c_{\max}$. We call "time upscaling" the possibility of using such a large time step, offered by an explicit precomputation of the propagator.

4.2.3 Complexity Analysis

In this section we derive the total complexity for the repeated squaring scheme in separated form (RS), as well as the subsequent upscaled timestepping (UTS). We will first formulate the total computational cost as a function of N—the initial data is on an N-by-N grid—as well as the various values of the ϵ-ranks r of submatrices corresponding to different wave vectors $(j', \mathbf{m}', \nu')$. In later sections we will carefully analyze how those ranks themselves depend on N and on the geometry of the speed of sound, $c(x)$.

We would like to make clear that the complexity estimates we are about to derive refer to the total number of operations when we fix some small threshold ϵ below which the singular values of submatrices are discarded. In practice, we observe that ϵ is very well correlated to the overall L^2 accuracy of the method (see Section 4.3), although we do not prove the connection on a rigorous level in this thesis. This observation of course assumes that the first time step in the initialization is itself made sufficiently accurate by taking Δt sufficiently small. Deriving accuracy estimates would imply dealing with sampling issues, namely that a function cannot be compactly supported both in space and in frequency. To obtain an L^2 accuracy $\overline{\epsilon}$ at time $T = 1$ we suspect that the threshold ϵ needs to depend on N and $\overline{\epsilon}$ like $O(\overline{\epsilon} N^{-1})$ as $N \to \infty$, in order to compensate for the cumulative error introduced by repeated squaring. In this scenario, the power of N in each complexity estimates would need to be incremented by an arbitrarily small number δ – at the expense of a constant depending on the choice of δ.

As always, we assume that $c(x)$ is C^∞. We measure complexity in terms of elementary floating point operations (flops). Let us remark once and for all that the two token indices ν and ν' do not play a role in the complexity analysis since we are interested in asymptotic results – up to constants.

4.2.3.1 Initialization

We start by observing that the initialization step of the repeated squaring can be done in $O(N^2 \log N)$ steps. Indeed, applying an inverse wavelet packet transform to find $\psi^j_{m_1,n_1}(x_1)$ takes $O(N \log N)$ operations. Performing multiplication by a function or differentiation takes at most the time of a FFT, $O(N \log N)$. Finally, applying a direct transform costs $O(N \log N)$ again. Since there are $O(N)$ values of the indices (j, m_1, n_1), and a constant ϵ-separation rank C_ϵ, the overall complexity is $O(C_\epsilon N^2 \log N)$.

As we will see, initialization happens to make for a negligible fraction of the total computing time.

4.2.3.2 Matrix Multiplication

Let us now consider the iteration step in the repeated squaring procedure, as described in Section 4.2.2.2. Fix a fine time step Δt and an upscaled time step $\tau = 2^{n^*}\Delta t$. We need to consider each scale j separately and recall that, by sparsity, we can always assume that j' is comparable to j. Ranks of submatrices are simply denoted by r, but let us keep in mind that they depend on the time step τ and the frequency indices $j, \mathbf{m}, j', \mathbf{m}'$, as well as on the desired accuracy level ϵ.

1. For fixed wave vectors, each subscript n_1 or n_2, or their counterpart with primes, takes on $O(2^j)$ values. Since both k and k' take on r values, a matrix such as $U^{\mathbf{k}}$ or $V^{\mathbf{k}}$ is of size $O(2^{2j}r^2)$. Each element of $U^{\mathbf{k}}$ or $V^{\mathbf{k}}$ takes $O(2^j)$ operations to compute, so the total complexity for forming two such matrices is $O(2^{3j}r^2)$.

2. One QR decomposition of $U^{\mathbf{k}}$ or $V^{\mathbf{k}}$ takes on $O(2^{2j}r^4)$ operations. Each middle factor $R_U \Sigma R_V^t$ is of size at most $O(r^2)$-by-$O(r^2)$, so the SVD of their product costs $O(r^6)$.

3. After performing the center SVD, only $O(r)$ singular values are kept (above the threshold ϵ), so we can trim both R_U and R_V to $O(r)$ columns. Computing each $U = Q_U R_U$ and $V = Q_V R_V$ takes $O(2^{2j}r^3)$ operations.

4. In the binary reduction using QR and SVDs, each matrix U has size $O(2^{2j})$-by-r, and let us grossly over-estimate r by its maximum $\max r$ over $(j', \mathbf{m}')$. Set M the total number of relevant indices $(j', \mathbf{m}')$. The total complexity for the binary reduction is $O(M(2^{2j}(\max r)^2 + (\max r)^3))$.

Steps 1, 2 and 3 above are to be repeated for each value of $(j', \mathbf{m}')$. In addition, there is an outer loop on the output wave vector $(j, \mathbf{m})$. The total complexity for one time doubling in the separated wave-atom repeated squaring algorithm is therefore

$$\text{Compl}(RS, \text{one step}) \leq C \cdot \sum_{j,\mathbf{m}} \sum_{j',\mathbf{m}'} (2^{3j} r^2 + 2^{2j} r^4 + r^6 + 2^{2j} (\max r)^2 + (\max r)^3).$$

For each j there are $O(2^{2j})$ different values of $\mathbf{m}$. We can use the inequality

$$\sum_{j',\mathbf{m}'} r^p \leq (\max r)^{p-1} \, (\sum_{j',\mathbf{m}'} r),$$

as well as the obvious $\max r \leq \sum_{j',\mathbf{m}'} r$, simplify and obtain

$$\text{Compl}(RS, \text{one step}) \leq C \cdot \sum_{j} \left[(2^{5j} \max r + 2^{4j} (\max r)^3 + 2^{2j} (\max r)^5) \, (\sum_{j',\mathbf{m}'} r) \right]. \quad (4.21)$$

The number of time doublings is small and depends remarkably little on the choices we make for Δt and τ. As long as both quantities are taken to depend inverse polynomially on N, the number of grid steps per dimension, the number of time doublings is $O(\log N)$. The total complexity for the repeated squaring is therefore $\log N$ times the right-hand side in equation (4.21).

We give very precise estimates for $\max_{j',\mathbf{m}'} r$ in Section 4.4 and $\sum_{j',\mathbf{m}'} r$ in Section 4.5, as a function of τ, ϵ and j (uniformly over $\mathbf{m}$). In Section 4.6, we improve the bounds of Sections 4.4 and 4.5 to take into account important special cases. In all cases, ranks and sums of ranks depend weakly on ϵ, namely they are all of order $O(\epsilon^{-1/M})$ for all $M > 0$ (with a constant depending on M.) This slow growth rate is the signature of spectral accuracy.

The simple choice $\tau = \frac{1}{\sqrt{N}}$, for example, is advantageous. We show in Theorem 4.1 that the worst-case estimate is $\sum r \simeq \max r \leq C_\epsilon \cdot 2^j$. In that case, the total complexity becomes

$$\text{Compl}(RS, \text{worst}) \leq C_\epsilon N^4 \log N.$$

For wave guides (when $c(x)$ depend only on x_1 or x_2, but not both,) the estimates become $\sum r \simeq \max r \leq C_\epsilon \cdot 2^{j/2}$, and

$$\text{Compl}(RS, \text{wave guide}) \leq C_\epsilon N^3 \log N.$$

Other choices for τ give rise to a variety of different complexity estimates. The situation is summarized in Figure 4.11.

4.2.3.3 Upscaled Timestepping

As we saw earlier, the complexity of the wave atom transform is $O(N^2 \log N)$.

The separated wave atom matrix should be unfolded into its classical form at every upscaled time step. For fixed wave vectors, the submatrix $E^{n_2,n_2'}_{n_1,n_1'}$ is of size $O(2^{2j})$-by-$O(2^{2j})$ and comes in separated form with rank r. Because of the restriction on nearby positions, all but $B_j = \max(2^j, \tau 2^{2j})$ rows and columns are kept, around $n_1' = n_1$ and $n_2' = n_2$. These rows and columns are easy to identify in the separated components as well, see equation (4.12), resulting in matrices u^k and v^k of size B_j-by-r. Explicitly forming $E^{n_2,n_2'}_{n_1,n_1'}$ from its separated components is a matrix-matrix product which takes $O(B_j^2 r)$ operations. The re-indexing of the relevant $O(B_j^2)$ elements into the classical form $E^{n_1',n_2'}_{n_1,n_2}$ takes $O(B_j^2)$ operations. Note that the latter matrix is band-diagonal with band $O(2^{-2j} B_j^2)$.

One matrix-vector product involving $E^{n_1',n_2'}_{n_1,n_2}$ then takes B_j^2 operations. Assuming that the solution $u(t)$ has a full set of wave atom coefficients (no particular sparsity pattern), then unfolding must be done for each wave vectors $(j', \mathbf{m}')$ (indexing columns) and $(j, \mathbf{m})$ (indexing rows), resulting in a total complexity

$$\text{Compl}(UTS, \text{one step}) \leq C \cdot \sum_{j,\mathbf{m}} B_j^2 (\sum_{j',\mathbf{m}'} r),$$

which can be rewritten more explicitly as

$$\text{Compl}(UTS, \text{one step}) \leq C \cdot \sum_j \left[(2^{4j} + \tau^2 2^{6j}) \, (\sum_{j',\mathbf{m}'} r) \right]. \tag{4.22}$$

Since T/τ upscaled time steps are necessary to reach time T (a multiple of τ), then the total complexity is the right-hand side of (4.22) multiplied by τ^{-1}.

For example, when τ is chosen as $\frac{1}{\sqrt{N}}$, then inspection of Theorems 4.1 and 4.4 reveals that the complexity estimate becomes

$$\text{Compl}(UTS, \text{worst case}) \leq C \cdot N^3$$

in the worst case, and

$$\text{Compl}(UTS, \text{wave guide}) \leq C \cdot N^{2.75}$$

in the case of wave guides. Estimates for different τ are summarized in Figure 4.11. By comparison, recall that a pseudospectral method would be $O(N^3 \log N)$.

Complexity and computational times can yet be improved when the wave atom expansion of the solution is uniformly sparse in time. Assume that $u(t)$, $0 \leq t \leq T$, can be approximated to accuracy ϵ in ℓ_2 using a fraction $\rho < 1$ of all wave numbers ξ_μ – not necessarily the same ones for different times. Then only the submatrices corresponding to those wave numbers must be computed at all, resulting in a direct net improvement of ρ of the complexity estimate for timestepping. For example, when $u(t)$ is a single bandlimited wavefront, then we can expect $\rho \simeq N^{-1/2}$. The corresponding total complexity becomes $O(N^{2.5})$ in general and $O(N^{2.25})$ for wave guides.

Complexity gains due to sparsity of the solution are harder to obtain for the repeated squaring, because it would demand identifying in advance which wave vectors are going to contribute in the yet unknown solution at dyadic times. These wave vectors are part of a "fat" manifold in phase-space. Such information could be obtained from a geometrical optics solver such as the phase-flow method [90], but we do not consider such a refinement in the present study.

Even though our complexity estimates may appear somewhat pessimistic, in particular for the core repeated squaring, it is worth keeping in mind that the result of the computation is not just one solution to the wave equation—*it is the whole Green's function in compressed form.* In particular, physical information of propagation of high frequencies can be read directly from the wave atom matrix representation.

4.3 Numerical Implementation and Examples

In this section, we apply the algorithm of Section 4.2.2 to several sample media. Theoretical studies of some of these representative media will be presented in Section 4.6. We used the orthonormal basis variation of Section 2.5 in all numerical experiments in this section.

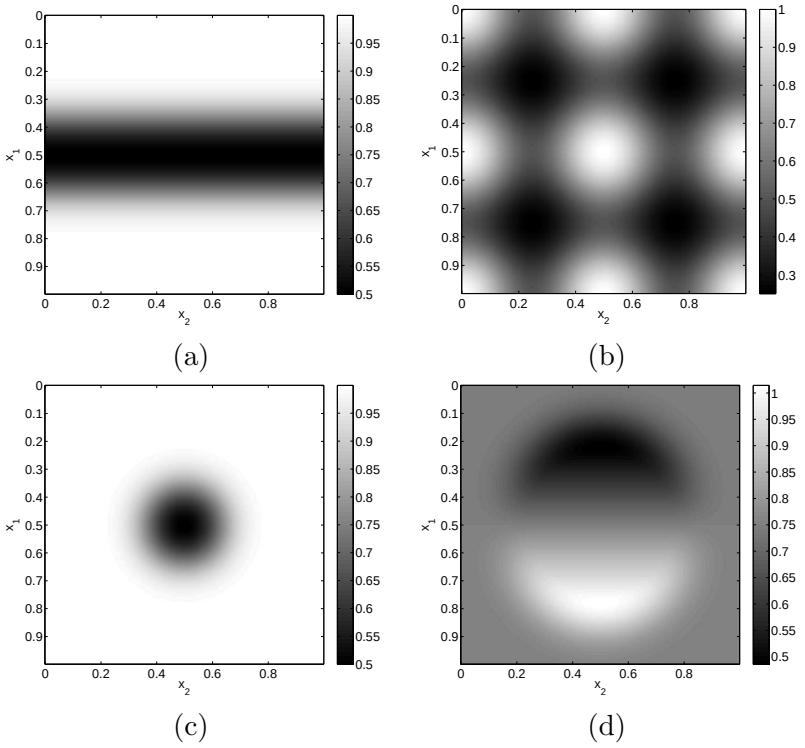

Figure 4.3: Four representative acoustic media. (a) wave guide, (b) bumps, (c) Gaussian converging lens, and (d) linear mirror.

We study four typical velocity fields defined over the unit square $[0,1)^2$:

- Wave guide (Figure 4.3(a)). The index of refraction is defined by

$$\eta(x_1, x_2) = 1 + \exp(-64 \times (x_1 - \frac{1}{2})^2).$$

- Bumps (Figure 4.3(b)). The wave speed is a simple trigonometric polynomial,

$$c(x_1, x_2) = \frac{(3 + \sin(4\pi x_1)) \cdot (3 + \sin(4\pi x_2))}{16}.$$

- Converging lens (Figure 4.3(c)). The index of refraction is given by

$$\eta(x_1, x_2) = 1 + \exp(-64 \times ((x_1 - \frac{1}{2})^2 + (x_2 - \frac{1}{2})^2).$$

- Linear mirror (Figure 4.3(d)),

$$c(x_1, x_2) = (0.75 + \rho(x_1, x_2)) \cdot (x_1 - \frac{1}{2})$$

  where ρ is a radial window function which smoothly extracts the center part of the unit square $[0, 1)^2$.

In each of these four cases, we apply the algorithm presented in Section 4.2.2 to generate the propagator $E(\tau)$ at time $\tau = 1/16$. The initial time step Δt used is set to 2^{-10}. The thresholding constant ϵ is chosen to be 10^{-4} and the grid size N is 128. As we pointed out already, the matrix $E(\tau)$ is organized as a collection of submatrices, which are indexed by row index $(j, \mathbf{m}, \nu)$ and column index $(j', \mathbf{m}', \nu')$. For each of the four media, the corresponding plot in Figure 4.4 describes the time dependence of the ϵ-rank. The solid curve is the maximum ϵ-rank over all submatrices, while the broken curve is the maximum of the sums of the ϵ-ranks over all column indices $(j', \mathbf{m}', \nu')$ (for a fixed row index $(j, \mathbf{m}, \nu)$). We compute these values at the dyadic time steps appeared in the construction $E(\tau)$, namely $t_n = 2^n \cdot \Delta t$, and linearly interpolate the value at other times.

We use two typical initial condtions to study our upscaled timestepping algorithm. The "line" initial condition (Figure 4.5(a)) is a smoothed indicator of $\{(x_1, x_2) : x_2 = \frac{1}{2}\}$ while the "pulse" (Figure 4.5(b)) is a smoothed delta function at the center of the domain. Both initial conditions, which have significant energy in the high frequency modes, are adequate for testing the numerical dispersion.

For each acoustic medium, we apply the upscaled time-stepping algorithm on these two initial conditions. We are particularly interested in conservation of the energy and accuracy of the wave profile. Since we start from the equation

$$p_{tt} - c^2(x)\Delta p = 0,$$

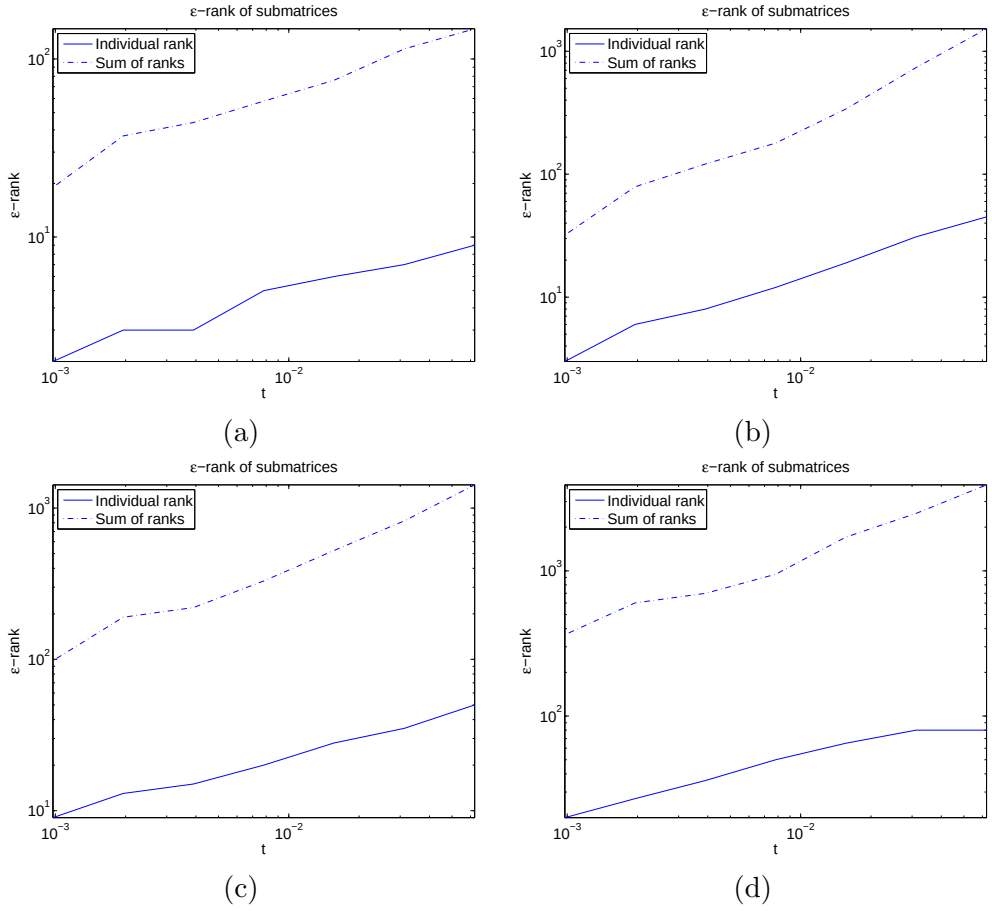

Figure 4.4: ϵ-rank of the submatrices. (a) wave guide, (b) bumps, (c) Gaussian converging lens, and (d) linear mirror.

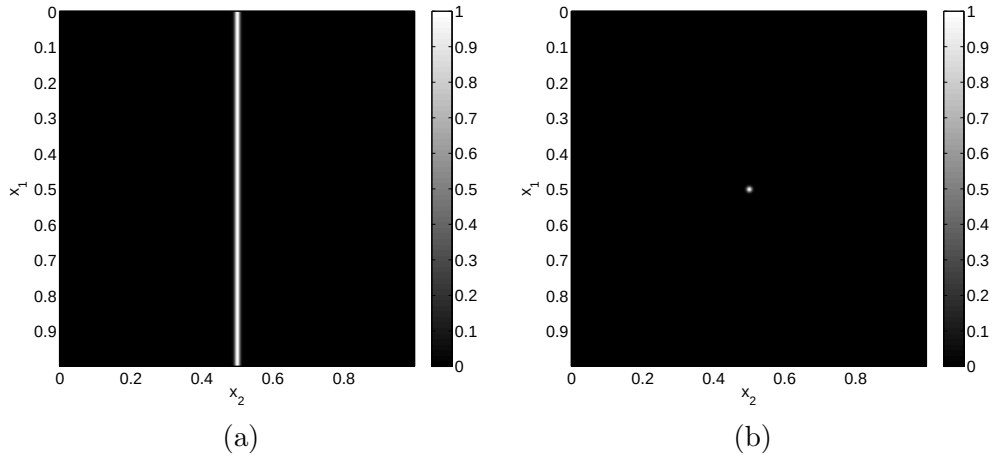

Figure 4.5: Initial condition used in the upscaled time-stepping algorithm.

the correct conserved energy is

$$\int \frac{|p_t|^2}{c^2(x)} + |\nabla p|^2 dx.$$

Figure 4.6 summarizes the time dependency of the relative errors of the energy integral (the solid curve) and the wave profile (the broken curve).

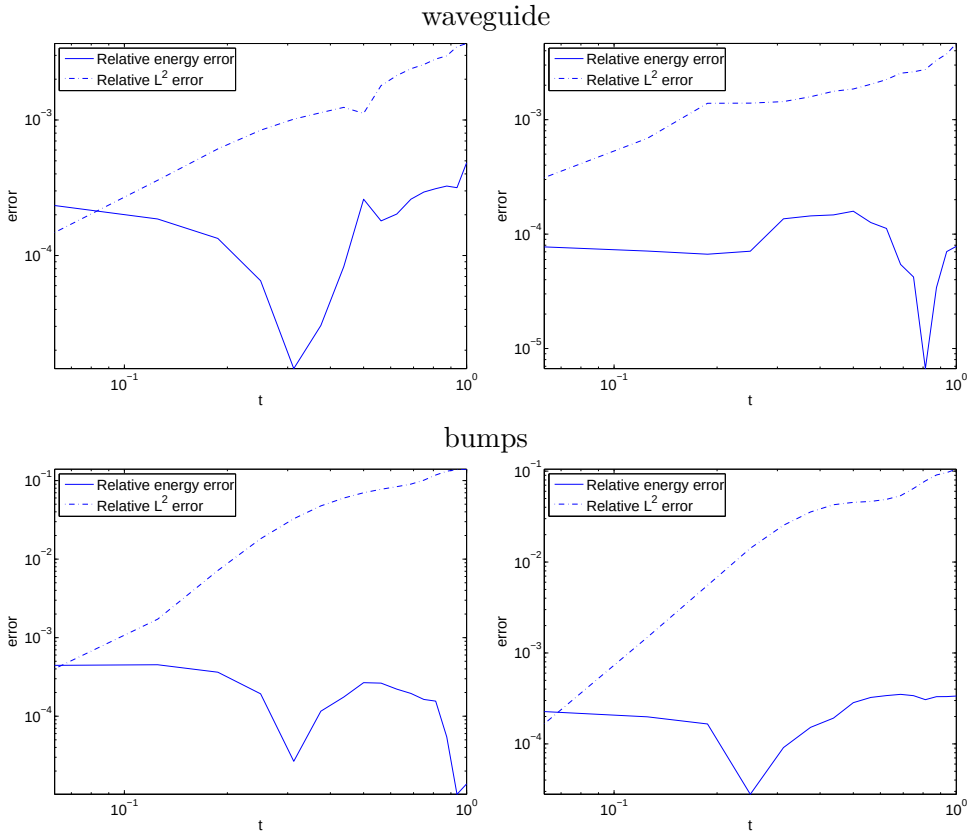

Figure 4.6: Relative error of the energy integral and the wave field. Left: "line" initial condition. Right: "pulse" initial condition.

It is well known that standard finite difference methods for hyperbolic equations often suffer from the problem of excessive numerical dispersion. This is particularly obvious when one uses a typical central-difference leapfrog scheme. In the following two experiments, we compare the numerical dispersion phenomenon in our upscaled time-stepping algorithm and the standard leapfrog algorithm. The time step and the grid size are chosen to be the same for both algorithms.

The first experiment involves the waveguide acoustic media and the "pulse" initial condition. The three images in Figure 4.8 show the solution at $t = 1/2$ and $t = 1$ computed using our method and the solution at $t = 1$ computed using the Leapfrog algorithm respec-

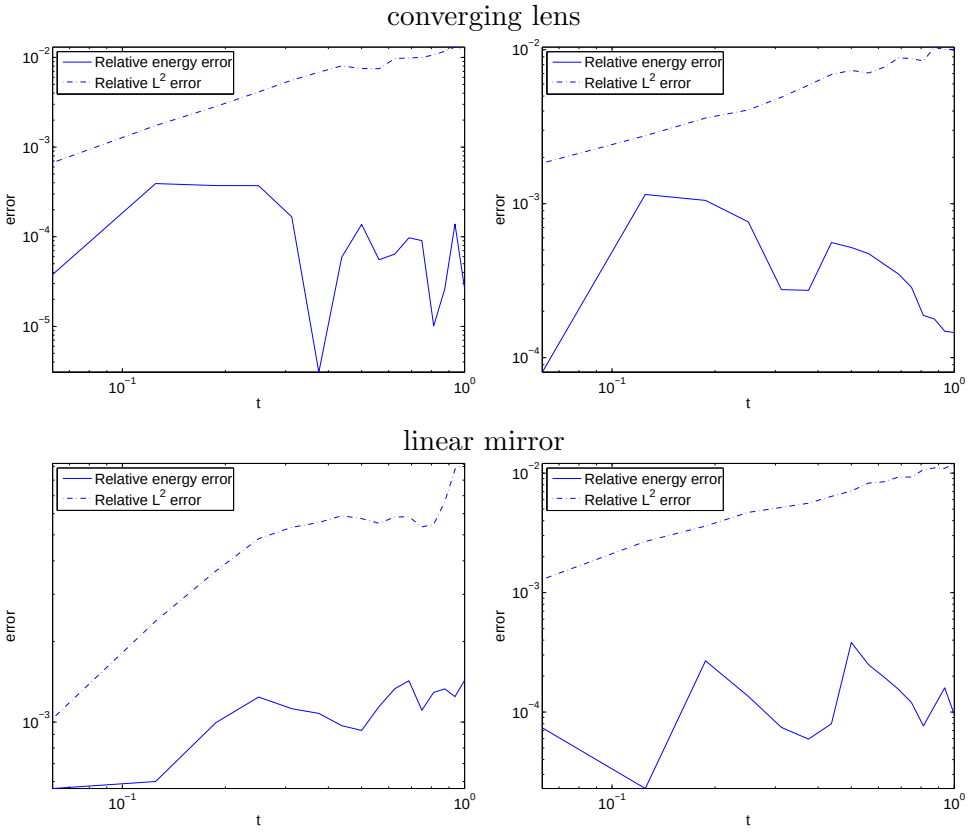

Figure 4.7: Relative error of the energy integral and the wave field. Left: "line" initial condition. Right: "pulse" initial condition.

tively. Notice that the ripples, which are the direct consequence of the numerical dispersion, are clearly observable in the leapfrog solution. The second experiment (Figure 4.9), which uses the bump acoustic media and the "line" initial condition, demonstrates the same phenomenon.

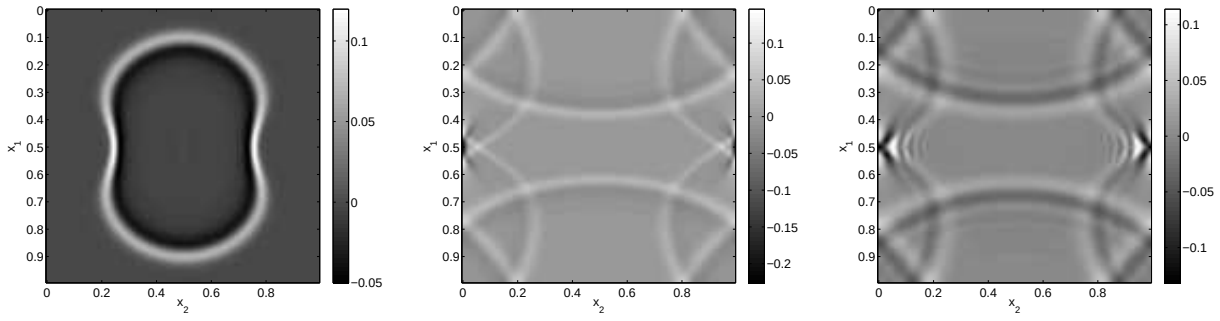

Figure 4.8: Numerical dispersion. Waveguide media and "pulse" initial condition. Left: $t = 1/2$, wave atom method. Middle: $t = 1$, wave atom method. Right: $t = 1$, finite difference method.

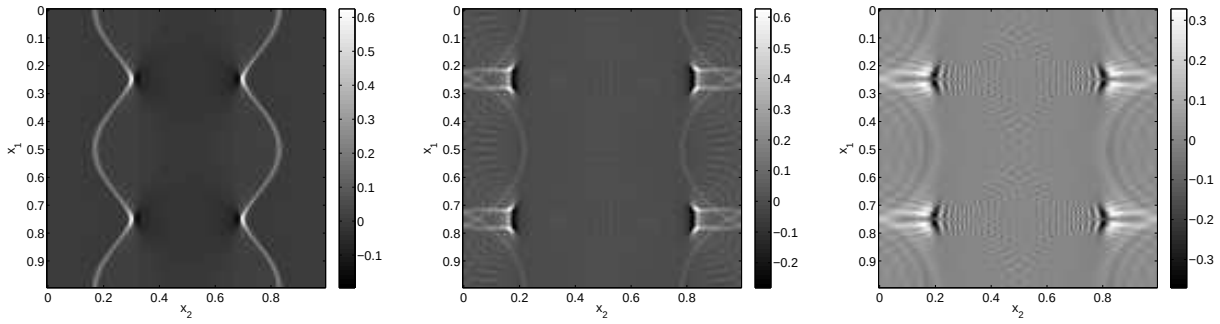

Figure 4.9: Numerical dispersion. Bump media and "line" initial condition. Left: $t = 1/2$, wave atom method. Middle: $t = 1$, wave atom method. Right: $t = 1$, finite difference method.

In the last experiment, we study the complexity of the upscaled time-stepping algorithm. As stated in Section 4.2.3, for certain types of acoustic media (e.g., wave guides), the upscaled timestepping algorithm has lower complexity compared to the standard spectral or pseudospectral methods, especially when the spatial discretization is refined. In fact, we are able to observe this fact even when N is relatively small. Figure 4.10 presents the time spent on applying a single upscaled time-step for various discretization size. For both the waveguide and bump media, the curve of the upscaled time-stepping algorithm grows much more slowly, and it becomes more efficient than the standard spectral method when N is larger than 128. This observation are in complete conformity with the complexity estimates in Section 4.2.3.

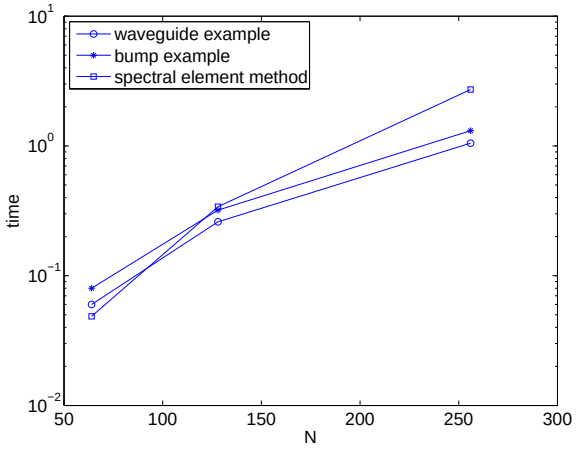

Figure 4.10: Computational time of a single upscaled time step. In all cases, the small time step is $\Delta t = 1/1024$.

4.4 Rank Estimates

The ϵ-rank r of a (possibly infinite) matrix A_{ij} is the smallest number r such that A_{ij} can be approximated up to accuracy ϵ by a matrix of rank r in ℓ^2,

$$\|A_{ij} - \sum_{k=1}^{r} u_i^k v_j^k\|_2 \le \epsilon.$$

The ϵ-separation rank, or just ϵ-rank r of a function $f(x_1, x_2)$ is the smallest number of separated components $u^k(x_1)v^k(x_2)$ necessary to approximate $f(x_1, x_2)$ up to accuracy ϵ in L^2, i.e.,

$$\|f(x_1, x_2) - \sum_{k=1}^{r} u^k(x_1)v^k(x_2)\|_2 \le \epsilon.$$

The main theoretical result of this chapter is a sharp bound on the ϵ-rank of reordered submatrices of the propagator in the wave atom frame. As detailed earlier each submatrix of interest has row index (n_1, n_2) vs. column index (n_1', n_2'), but the separation isolates (n_1, n_1') vs. (n_2, n_2'). Hence the necessity of *reordering* the entries, to prepare the submatrix for standard low-rank approximation. Notice that the size of the remainder, no more than ϵ, is however measured in ℓ^2 in the *original* form (n_1, n_2) vs. (n_1', n_2'), in complete conformity with the goal of bounding the overall ℓ^2 norm of the error on the propagator.

Theorem 4.1. *Assume the velocity profile $c(x)$ is C^∞. Consider the submatrix $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$ obtained by fixing $j, \mathbf{m}, \nu$ and $j', \mathbf{m}', \nu'$ in the wave atom representation of the propagator $E(t)$. It is of size $O(2^{2j})$-by-$O(2^{2j})$, where $|\xi_\mu| = O(2^{2j})$ and at finest scale $2^{2j} \simeq N$. After reordering $(n_1, n_2; n_1', n_2') \to (n_1, n_1'; n_2, n_2')$, $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$ has ϵ-rank r bounded as follows.*

- *for $t \lesssim 2^{-j}$, $r \le C_\epsilon \cdot (1 + t2^{2j})$,*
- *for $2^{-j} \lesssim t \lesssim 2^{-j/2}$, $r \le C_\epsilon \cdot 2^j$,*
- *for $2^{-j/2} \lesssim t \le T$, $r \le C_\epsilon \cdot t^2 2^{2j}$,*

with $C_\epsilon \le C_M \epsilon^{-1/M}$, for all $M > 0$, and C_ϵ also depends on T.

The various values taken on by the bound on r are summarized in Figure 4.11. Notice that for large times the rank r is always obviously bounded by $C2^{2j}$, the size of the submatrix.

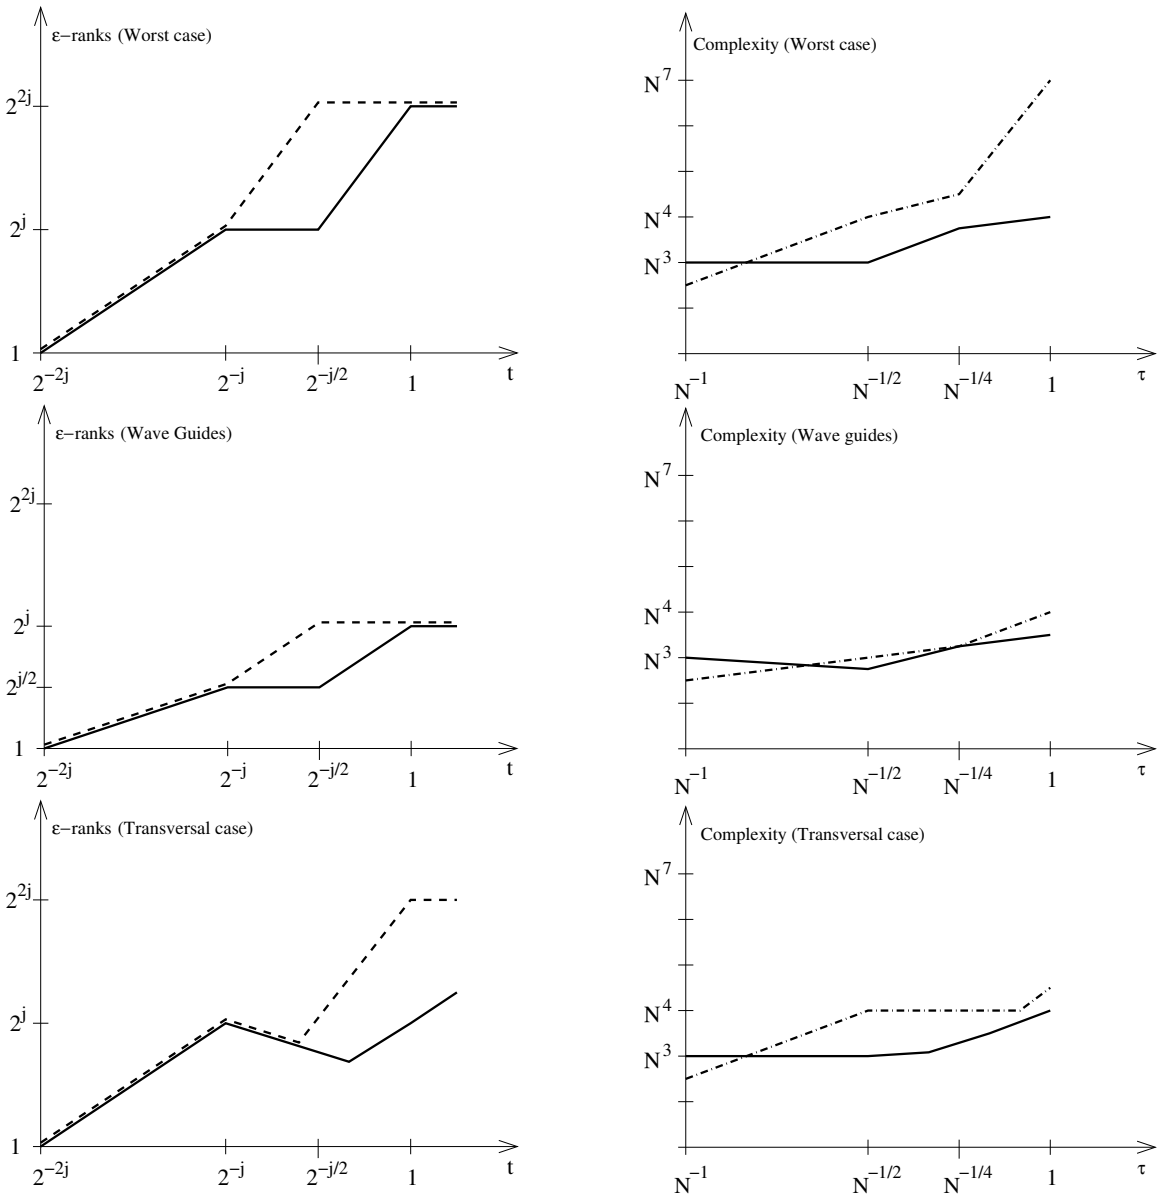

Figure 4.11: Left: schematic illustration of the bounds on the ϵ-ranks of submatrices. Solid line: bound on individual ranks, as in Theorem 4.1. Dashed line: bound on sums of ranks over wave vectors, as in Corollary 4.1. Right: schematic illustration of the overall complexity. Solid line: upscaled time-stepping. Dash-dotted line: repeated squaring. All values of time, ranks and complexity should be understood modulo multiplicative constants and even $\log N$ factors. The portion of the t-axis $2^{-j} \lesssim \tau \lesssim 2^{-j/2}$ corresponds to the region of interest for numerical computations.

Before proving this result, we need to recall that the propagator $E(t)$ can be approximated by an oscillatory integral, called the Lax parametrix, of the form

$$E(t)u(x) = \sum_{\ell=\pm} \int e^{i\Phi_\ell(x,\xi,t)} a_\ell(x,\xi,t)\hat{u}(\xi)\, d\xi + R_1(t)u(x).$$

This formula is only valid for small times $0 \leq t < T$ before caustics. For smooth C^∞ media, each phase function is C^∞, positive-homogeneous of degree 1 in ξ, and solves a Hamilton-Jacobi equation,

$$\frac{\partial \Phi_\pm(x,\xi,t)}{\partial t} = \pm c(x)|\nabla_x \Phi_\pm(x,\xi,t)|, \qquad \Phi_\pm(x,\xi,0) = x \cdot \xi. \tag{4.23}$$

Each matrix-valued amplitude a_ℓ is a symbol of order 0 and type $(1,0)$, i.e., componentwise,

$$|\partial_\xi^\alpha \partial_x^\beta a_\ell(x,\xi,t)| \leq C_{\alpha,\beta,t}(1+|\xi|)^{-|\alpha|}, \qquad \text{for all } \alpha.$$

This condition is denoted $a_\ell \in S^0$. The remainder $R_1(t)$ is smoothing in the sense that it turns tempered distribution into C^∞ functions. This is the same setting as in [15], to which we refer for details and justifications.

In what follows, we consider $x \in \mathbb{R}^2$, but $u(t,x)$ with support in a subset Ω inside the open unit square $]0,1[^2$. Without loss of generality, we can make $a_\ell(x,\xi,t)$ compactly supported through multiplication by an adequate cutoff equal to one on Ω and tapering smoothly to zero outside $[0,1]^2$.

We will need the following two simple lemmas in the sequel.

Lemma 4.2. *Let $f \in C^\infty(]0,1[^2)$. Then the ϵ-separation rank of f obeys $r_\epsilon(f) \leq C_M \epsilon^{-1/M}$ for all $M > 0$, and some constant $C_M > 0$. Furthermore, for each positive integer $s > 0$ there exists a constant C_s such that the same rank-$r_\epsilon(f)$ decomposition has error $C_s\epsilon$ when measured in the Sobolev space $W^{s,p}$, for all $1 \leq p \leq \infty$.*

Proof. It suffices to notice that the Fourier series coefficients of f decay like

$$|\hat{f}[k]| \leq C_M |k|^{-M}.$$

Each Fourier mode is separable. Truncating the Fourier series to $|k| \leq K$, by means of

$O(K^2)$ terms, results in a squared L^2 error

$$\sum_{|k|>K} |\hat{f}[k]|^2 \leq C_M^{(0)} K^{-M}.$$

for all $M > 0$, and some other constant $C_M^{(0)} > 0$. Hence K can be chosen less than $C_M \epsilon^{-1/M}$ for some adequate choice of constant C_M.

The expression of the square of the error in $W^{s,2}$ is (up to constants)

$$\sum_{|k|>K} |k|^{2s} |\hat{f}[k]|^2 \leq C_M^{(s)} K^{-M},$$

for all $M > 0$ but where the constant $C_M^{(s)}$ is likely larger than $C_M^{(0)}$. Since $K \leq C_M \epsilon^{-1/M}$, we certainly have a $W^{s,2}$ error bounded by $C_s \epsilon$ for some constant C_s depending on s.

When $p \neq 2$, we conclude using the continuous Sobolev inclusion $W^{2+s,2} \subset W^{s,p}$, valid for all $1 \leq p \leq \infty$ in two dimensions.

□

Lemma 4.3. *Let $\Phi(x, \xi, t)$ solve either Hamilton-Jacobi equation (4.23). Then, for $0 \leq t \leq T$ while the equation is well posed, the Hessian obeys*

$$\nabla_x \nabla_x \Phi(x, \xi, t) = t\psi(x, \xi, t),$$

where each component of ψ is C^∞ away from $\xi = 0$, and positive-homogeneous of degree 1 in ξ.

Proof. Write $\nabla_x \nabla_x \Phi(x, \xi, t)$ as the integral $\pm \int_0^t \nabla_x \nabla_x c(x) |\nabla_x \Phi(x, \xi, s)| \, ds$, where the integrand has the same smoothness and homogeneity properties as Φ itself. □

Proof of theorem 4.1. Let $\epsilon > 0$. We seek a bound on the number r of separated terms in

$$\langle \varphi_\mu \mathbf{e}_\nu, E(t) \varphi_{\mu'} \mathbf{e}_{\nu'} \rangle = \sum_{k=1}^{r} u^k_{n_1, n_1'}(t) v^k_{n_2, n_2'}(t) + R_{\mathbf{n}, \mathbf{n}'}(t),$$

where $R_{\mathbf{n},\mathbf{n}'}(t)$, as a matrix with row subscript $\mathbf{n}$ and column subscript $\mathbf{n}'$, has ℓ^2 norm less than ϵ. Note that all the quantities r, u, v and R depend on the parameters $j, \mathbf{m}, \nu$ and

$j', \mathbf{m}', \nu'$, but we drop this dependence for simplicity of notations.

By construction, wave atoms have rank 2 in the frequency domain, namely

$$\hat{\varphi}_\mu(\xi) = \left[\hat{\psi}^j_{m_1,+}(\xi_1)\hat{\psi}^j_{m_2,+}(\xi_2) + \hat{\psi}^j_{m_1,-}(\xi_1)\hat{\psi}^j_{m_2,-}(\xi_2) \right] e^{-i2^{-j}n_1\xi_1} e^{-i2^{-j}n_2\xi_2},$$

Without loss of generality, at the expense of at most quadrupling the constants in front of each estimate, we will drop the second term in the above parenthesis. This results in considering each wave atom as having only one separated bump in the frequency plane, i.e, having rank 1. We keep the notation $\hat{\varphi}_\mu(\xi)$ for these "amputated" atoms.

Call S_μ the support of $\hat{\varphi}_\mu(\xi)$; it can be inscribed in a ball centered at ξ_μ, and of radius equal to $2^{j+1}\sqrt{2}\pi$. We will denote by $\chi_\mu(\xi)$ a smooth and separable indicator function, equal to one on S_μ, and zero on the complement of the larger set $S_\mu + \{\xi : |\xi| \le 2^j\}$.

Denote by $E(\xi, \eta, t)$ the frequency kernel of $E(t)$, namely

$$\widehat{E(t)u}(\xi) = \int E(\xi, \eta, t)\hat{u}(\eta)\, d\eta.$$

By Parseval, the matrix elements are

$$\langle \varphi_\mu, E(t)\varphi_{\mu'} \rangle = \int\int \hat{\psi}^j_{m_1,+}(\xi_1)\hat{\psi}^j_{m_2,+}(\xi_2) e^{-i2^{-j}n_1\xi_1} e^{-i2^{-j}n_2\xi_2} E(\xi, \eta, t)$$

$$\times \overline{\hat{\psi}^{j'}_{m_1',+}(\eta_1)\hat{\psi}^{j'}_{m_2',+}(\eta_2)} e^{i2^{-j'}n_1'\eta_1} e^{-i2^{-j'}n_2'\eta_2}\, d\xi_1 d\xi_2 d\eta_1 d\eta_2, \tag{4.24}$$

where the kernel is

$$E(\xi, \eta, t) = K(\xi, \eta, t) + R_1(\xi, \eta, t), \tag{4.25}$$

where

$$K(\xi, \eta, t) = \sum_{\ell=\pm} \int e^{i(\Phi_\ell(x,\eta,t) - x\cdot\xi)} a_\ell(x, \eta, t)\, dx. \tag{4.26}$$

In what follows we will drop the sum and the subscript ℓ (at the expense of doubling the separation rank,) because $\ell = -$ is totally analogous to $\ell = +$. We now seek results of separation of $K(\xi, \eta, t)$ in both ξ and η, on the frequency support of each wave atom, i.e.,

$$\chi_\mu(\xi)\chi_{\mu'}(\eta)K(\xi, \eta, t) = \sum_{k=1}^{r} K_k^{(1)}(\xi_1, \eta_1, t)K_k^{(2)}(\xi_2, \eta_2, t) + R_2(\xi, \eta, t). \tag{4.27}$$

The following lemma shows that the size of the remainder R_2 directly translates into a remainder of comparable size for the submatrix of interest.

Lemma 4.4. *Let $T(\xi,\eta)$ be any kernel defining by extension a bounded operator T on $L^2(\mathbb{R}^2)$. As usual, we denote $\mu = (j,\mathbf{m},\mathbf{n})$ and $\boldsymbol{\varphi}_{\mu\nu} = \varphi_\mu \mathbf{e}_\nu$. For any $(j,\mathbf{m},\nu)$ and $(j',\mathbf{m}',\nu')$,*

$$\| \langle \boldsymbol{\varphi}_{\mu\nu}, T\boldsymbol{\varphi}_{\mu'\nu'} \rangle \|_{\ell^2_{\mathbf{n}'}\to\ell^2_{\mathbf{n}}} \le \|T\|_{L^2_{\eta\nu'}\to L^2_{\xi\nu}} \tag{4.28}$$

Proof. By the tight frame property,

$$\|T\|_{L^2_{\eta\nu'}\to L^2_{\xi\nu}} = \| \langle \boldsymbol{\varphi}_{\mu\nu}, T\boldsymbol{\varphi}_{\mu'\nu'} \rangle \|_{\ell^2_{\mu'\nu'}\to\ell^2_{\mu\nu}}.$$

Choose any set of indices $(j,\mathbf{m},\nu)$ and $(j',\mathbf{m}',\nu')$. The conclusion follows by restricting the wave atoms matrix of T to rows indexed by $(j,\mathbf{m},\nu)$ and columns indexed by $(j',\mathbf{m}',\nu')$. □

Notice that the remainder R_1 does not pose any difficulty. Since it corresponds to a smoothing operator on a bounded domain, we have the bound

$$|R_1(\xi,\eta,t)| \le C_M(1+|\xi|+|\eta|)^{-M}, \qquad \text{for all } M>0,$$

so, in the spirit of Lemmas 4.2 and 4.4, a constant number $C_\epsilon \sim \epsilon^{-1/M}$ of (separable) Fourier modes suffices to approximate the submatrix coming from R_1 to accuracy $\epsilon/2$ in $\ell^2_{\mathbf{n}'}\to\ell^2_{\mathbf{n}}$.

It is also important to notice that for each fixed wave vector ξ_μ, only a few wave vectors $\xi_{\mu'}$ give rise to nonnegligible matrix elements. This is due to sparsity, and quantified in Section 4.5. For the time being we only need to observe that, for those nonnegligible entries, the wave vectors are comparable; very conservatively, $|j-j'| \le$ const. Also, the particular value of $\mathbf{m}$ will be seen not to play any significant role. As a consequence, ϵ-separation ranks essentially only depend on one of the two numbers j and j', say j. In the sequel, we will look for a bound on r which depends solely on j, understanding that it holds uniformly over all $j',\mathbf{m}'$ and $\mathbf{m}$.

We are now ready to split the proof into three parts, corresponding respectively to (1) coarse scales, (2) fine scales in the regions of nonstationary phase, and (3) fine scales near the locus of stationary phase.

4.4.1 Coarse Scales

The case of coarse scales, i.e., say $j = j' = 0$, needs to be considered separately because the phase $\Phi(x, \xi, t)$ has in general a kink at the origin in ξ, that is, a discontinuity in the gradient.

Let $g(x, \eta, t) = e^{i\Phi(x,\eta,t)} a(x, \eta, t)\chi_\mu(\eta)$, so that

$$K(\xi, \eta, t) = \int e^{-ix\cdot\xi} g(x, \eta, t)\, dx = \hat{g}(\xi, \eta, t),$$

where the Fourier transform is taken over the first variable only.

Take $\{\psi_\lambda\}$ a 2D separable wavelet orthonormal basis, with super-algebraic decay in both space and frequency, and expand $g(x, \cdot, t)$:

$$g(x, \eta, t) = (\sum_{\lambda\in\Lambda_1} + \sum_{\lambda\in\Lambda_2})\, c_\lambda(x, t)\psi_\lambda(\eta).$$

Determine the subset of subscripts Λ_2 such that

$$\sup_x \sup_{0\le t\le T} \sum_{\lambda\in\Lambda_2} |c_\lambda(x, t)|^2 \le \frac{\epsilon^2}{16\pi^2}.$$

Since $\nabla_\eta g(x, \cdot, t)$ is discontinuous at the origin, but otherwise C^∞ and compactly supported in a $O(1)$ region, it is a classical result from wavelet analysis that

$$|\Lambda_1| \le C_M \epsilon^{-1/M}, \qquad \text{for all } M > 0.$$

This constant number of important subscripts in Λ_1 correspond to large scales as well as locations near the singularity.

Therefore,

$$K(\xi, \eta, t)\chi_\mu(\eta) = \sum_{\lambda\in\Lambda_1} \hat{c}_\lambda(\xi, t)\psi_\lambda(\eta) + R_2(\xi, \eta, t).$$

Each coefficient $c_\lambda(x, t)$ inherits the C^∞ smoothness of g, and is essentially supported near the unit cube. By lemma 4.2, the ϵ-separation rank of $\hat{c}_\lambda(\xi, t)$ is therefore $O(\epsilon^{-1/M})$ for all $M > 0$. In addition, each $\psi_\lambda(\eta)$ is separable, and the sum runs over at most $O(\epsilon^{-1/M})$ terms. So the overall separation rank for the sum is $O(\epsilon^{-1/M})$ as well.

Let us now check that $R_2(\xi, \eta, t)$ generates an error which is the correct fraction of ϵ in

L^2. The squared Hilbert-Schmidt norm of R_2, which bounds the squared L^2 norm, is

$$\begin{aligned}\int\int |R_2(\xi,\eta,t)|^2\, d\xi d\eta &= \int \sum_{\lambda\in\Lambda_2} |\hat{c}_\lambda(\xi,t)|^2\, d\xi \quad \text{by Plancherel-wavelets,}\\ &= \sum_{\lambda\in\Lambda_2} \int |c_\lambda(x,t)|^2\, dx \quad \text{by Plancherel-Fourier,}\\ &\leq \sup_{x\in[0,1]^2} \sum_{\lambda\in\Lambda_2} |c_\lambda(x,t)|^2 \leq \frac{\epsilon^2}{16\pi^2}.\end{aligned}$$

We can now apply Lemma 4.4 to obtain a remainder of size $\epsilon/2$ for each wave atom submatrix of R_2, with indices $\mathbf{n}, \mathbf{n}'$. Together with R_1's submatrices, also of size $\epsilon/2$, the overall remainder is of size at most ϵ. This finishes the proof for the coarse scales, with the result that $r = O(\epsilon^{-1/M})$.

4.4.2 Fine Scales, Stationary Phase

Consider μ and μ' such that $\chi_\mu(\xi) = \chi_{\mu'}(\xi) = 0$ in a neighborhood of the origin $\xi = 0$. We expect the integrand in equation (4.26) to be large near the points of stationary phase, i.e. $\xi = \nabla_x\Phi(x,\eta,t)$. For each $\delta > 0$, consider the sets

$$X_\xi^\eta(\delta) = \{x \in [0,1]^2; |\xi - \nabla_x\Phi(x,\eta,t)| \leq 2^j\delta\},$$

and their union

$$X_\mu^{\mu'}(\delta) = \{x \in [0,1]^2; \text{there exist } \xi \in \text{supp}\chi_\mu \text{ and } \eta \in \text{supp}\chi_{\mu'},\ |\xi - \nabla_x\Phi(x,\eta,t)| \leq 2^j\delta\}. \tag{4.29}$$

Our aim is to find a smooth indicator $p(x)$ equal to one for $x \in X_\mu^{\mu'}(\delta)$, and for which the restricted kernel

$$K_{\text{nonstat}}(\xi,\eta,t) = \int (1-p(x)) e^{i\Phi(x,\eta,t) - ix\cdot\xi} a(x,\eta,t)\, dx$$

is negligible in the L^2 sense, $\|K_{\text{nonstat}}\|_2 \leq \epsilon/4$. We will see in the next section that such an estimate holds provided δ is chosen large enough; let us accept for the moment that it can be taken of the form $\delta = O(\epsilon^{-1/M})$.

In this section we show how to build $p(x)$ as a sum of functions $q_k(x)$, which define

kernels

$$K_k(\xi,\eta,t) = \int q_k(x) e^{i\Phi(x,\eta,t) - ix\cdot\xi} a(x,\eta,t)\, dx, \tag{4.30}$$

such that each $K_k(\xi,\eta,t)\chi_\mu(\xi)\chi_{\mu'}(\eta)$ has ϵ-separation rank of order $O(\epsilon^{-1/M})$ in ξ and η, for all $M > 0$. An estimate on the overall rank is then expected, for then

$$K = K_{\text{nonstat}} + \sum_{k=1}^{N_B} K_k \tag{4.31}$$

will be well separated by $O(N_B\epsilon^{-1/M})$ terms for all $M > 0$. In the rest of this section we intend to estimate N_B as a function of t as well as justify smallness of the non-separated remainder.

The first observation is that the union in the definition of $X_\mu^{\mu'}(\delta)$ is not essential. More precisely, let σ be the Lyapunov exponent of the bicharacteristic Hamiltonian system,

$$\sigma = \sup_{t\geq 0} \frac{1}{t} \log \left(\sup_{x\in[0,1]^2} \sup_{\xi\in\mathbb{R}^2} |\nabla_x \nabla_\xi \Phi(x,\xi,t)| \right).$$

For any $\xi_0, \xi \in \text{supp}\chi_\mu$ and $\eta_0, \eta \in \text{supp}\chi_{\mu'}$, we have the estimates

$$|\xi_0 - \xi| \leq C2^j \qquad \text{and} \qquad |\eta_0 - \eta| \leq C'2^{j'} \leq C2^j.$$

A Taylor expansion of Φ around η_0 then reveals

$$X_\mu^{\mu'}(\delta) \subset X_{\xi_0}^{\eta_0}(Ce^{\sigma t}\delta). \tag{4.32}$$

This observation is important because it shows that the condition $|\xi - \nabla_x\Phi(x,\eta,t)| \leq 2^j\delta$ is the strongest definition of the neighborhood of the locus of stationary phase which still makes it independent of ξ and η.

The next step is to linearize the phase $\Phi(x,\eta,t)$ in η near some point $\eta_0 \in \text{supp}\chi_{\mu'}$. The whole point of partitioning the frequency plane into indicators of radius $O(2^j)$, when $|\eta| \sim 2^{2j}$, is precisely to make the remainder non-oscillatory. More precisely, for $\eta \in \text{supp}\chi_{\mu'}$, homogeneity of degree one in η implies the estimate

$$\partial_\eta^\alpha \left[\Phi(x,\eta,t) - \eta\cdot\nabla_\eta\Phi(x,\eta_0,t)\right] = O(|\eta|^{-|\alpha|/2}).$$

For a proof, see [15], appendix B (p.55); or [79], chapter IX, pp. 406–407. This nonlinear remainder can be absorbed in the amplitude, which we still denote $a(x,\eta,t)$ for simplicity,

$$a(x,\eta,t)\chi_{\mu'}(\eta) := e^{i(\Phi(x,\eta,t)-\eta\cdot\nabla_\eta\Phi(x,\eta_0,t))}a(x,\eta,t)\chi_{\mu'}(\eta),$$

without essentially changing its properties: the new amplitude $a\chi_{\mu'}$ is still of order zero and type $(1/2,0)$, i.e.,

$$|\partial_\eta^\alpha\partial_x^\beta a(x,\eta,t)\chi_{\mu'}(\eta)| \leq C_{\alpha,\beta}(1+|\eta|)^{-|\alpha|/2}. \tag{4.33}$$

The central argument now consists in performing Taylor expansions of the (linearized) phase in x within adequately small balls $B_{x_k}(\rho_k)$. Call $f(x) = \eta\cdot\nabla_\eta\Phi(x,\eta_0,t)$. Then

$$f(x) = f(x_k) + (x-x_k)\nabla f(x_k) + \frac{1}{2}(x-x_k)^t\nabla\nabla f(y)(x-x_k), \tag{4.34}$$

where $x, y \in B_{x_k}(\rho_k)$, and $\nabla\nabla f$ denotes the Hessian. The first genuinely non-separable contribution comes from the off-diagonal quadratic term x_1x_2. We can still have control over this term if we make it nonoscillatory, i.e., if we take ρ_k small enough that

$$\rho_k^2|\nabla\nabla f(x)| \leq C \qquad \text{for } x \in B_{x_k}(\rho_k). \tag{4.35}$$

The point is that the constant C is independent of j. The quadratic term can then be absorbed in the amplitude without essentially changing the latter, as was done previously for the linearization in η.

We are then led to the geometric problem of covering the set $X_\mu^{\mu'}(\delta)$ with the smallest possible number of balls $B_{x_k}(\rho_k)$ in which the quadratic term is non-oscillatory. Let us first lighten notations by writing $g(x)$ for either $\frac{\partial}{\partial x_1}\nabla_\eta\Phi(x,\eta_0,t)$ or $\frac{\partial}{\partial x_2}\nabla_\eta\Phi(x,\eta_0,t)$. Uniform boundedness of the quadratic term, as above, can be expressed as

$$|\nabla g(x)| \leq C\cdot 2^{-2j}\rho_k^{-2}. \tag{4.36}$$

As we saw in equation (4.32), the condition $x \in X_\mu^{\mu'}(\delta)$ can be reduced to $x \in X_\xi^{\eta_0}(C\cdot\delta)$, which in turn reads

$$|g(x)| \leq C\cdot 2^{-j}. \tag{4.37}$$

Notice that $g(x)$, like the sound speed $c(x)$, is C^∞ for times $t < T$ before breakdown of the Hamilton-Jacobi equation on plane wave initial conditions. We then claim that, for any smooth $g(x)$ (C^2 will suffice), the set where (4.37) holds can be covered by $N_B = O(2^j)$ balls in which (4.36) holds. The construction of such a covering necessarily depends on $g(x)$ itself, so we apologize to the reader for the following argument being a bit technical.

We switch to a continuous description of the problem by introducing a local *ball radius density* $\rho(x, j)$ which will help determine $\rho_k = \rho(x_k, j)$ at a collection of points x_k still to be determined. We set

$$\rho(x, j) = \frac{1}{\sqrt{2^j + 2^{2j}|\nabla g(x)|}}. \tag{4.38}$$

Two basic properties motivate this formula, namely that

- $|\nabla g(x_k)| \leq 2^{-2j}\rho(x_k, j)^{-2}$, as required, and
- $C \cdot 2^{-j} \leq \rho(x, j) \leq 2^{-j/2}$, for all $x \in [0, 1]^2$.

It is important for what follows to check that formula (4.38) is consistent as a definition of local radius, in the sense that

$$\sup_{x \in B_{x_k}(\rho_k)} \rho(x, j) \leq C_{c.o.} \cdot \rho(x_k, j). \tag{4.39}$$

This result is an easy consequence of Landau's inequality and is justified in the appendix. We call it the *constant overlap property*.

The collection of ball centers x_k is now determined as follows. Start from a Cartesian lattice $y_k = (k_1, k_2)b2^{-j}$ with k_1, k_2 integers and some small $b > 0$ to be determined. Assign a ball of center $\tilde{\rho}_k = \rho_k/5 = \rho(x_k, j)/5$ to y_k. The constant b is taken so that the union of all the balls $B_{y_k}(\tilde{\rho}_k)$ covers $[0, 1]^2$. In general the balls significantly overlap and the covering needs pruning, for instance by means of the following elementary covering lemma.

Lemma 4.5. *Let $\mathcal{G} = \{B_{y_k}(\tilde{\rho}_k)\}$ be a family of closed balls with uniformly bounded radius. Then there is a subfamily $\mathcal{F} \subset \mathcal{G}$ of pairwise disjoint balls such that*

$$\bigcup_{B_{y_k}(\tilde{\rho}_k) \in \mathcal{G}} B_{y_k}(\tilde{\rho}_k) \subset \bigcup_{B_{y_k}(\tilde{\rho}_k) \in \mathcal{F}} B_{y_k}(5\tilde{\rho}_k).$$

Proof. See [92], p. 7. □

The collection x_k then emerges as the centers of the remaining balls and the radii are chosen as $\rho_k = 5\tilde{\rho}_k$.

Notice that, by construction, each point in the unit square is covered by at most a constant number of balls $B_{x_k}(\rho_k)$ (independent of j or ρ_k). This is because the constant overlap property (4.39) can be iterated to yield

$$\sup_{x \in B_{x_k}(2\rho_k)} \rho(x, j) \asymp \rho(x_k, j).$$

(The notation $A \asymp B$ means $A \leq C \cdot B$ and $B \leq C \cdot A$ for some positive C which may depend on some parameters, depending on context.) The balls overlapping with $B_{x_k}(\rho_k)$ therefore have radius comparable to ρ_k, so there can only be a constant number of them.

We are now ready to estimate the number N_B of balls which cover $X \equiv X_\mu^{\mu'}(\delta)$. To every lattice point y_k, assign a weight

$$w_k = 2^{-2j} \sum_{x_\ell \in X : y_k \in B_{x_\ell}(\rho_\ell)} \frac{1}{\rho_\ell^2}. \tag{4.40}$$

Since there are $O(2^{2j}\rho_\ell^2)$ grid points y_k inside the ball $B_{x_\ell}(\rho_\ell)$, it is straightforward to check that

$$N_B \leq C \cdot \sum_k w_k.$$

On the other hand, the constant overlap property (equation (4.39)) entitles us to see $\sum_k w_k$ as a Riemann sum and bound

$$\sum_k w_k \leq C \cdot \int_X \frac{1}{\rho^2(x, j)}\, dx.$$

Using the definition (4.38), we get

$$N_B \leq C \cdot (2^j + 2^{2j} \int_X |\nabla g(x)|\, dx).$$

We claim that $\int_X |\nabla g(x)|\, dx \leq C \cdot 2^{-j}$. This fact follows from the following lemma, which is a simple reformulation of the co-area formula for BV functions. For our application, we let $\tilde{\epsilon} = 2^{-j}$.

Lemma 4.6. *Let $g \in C^2([0,1]^2)$. For all $\tilde{\epsilon} > 0$, let $X_{\tilde{\epsilon}} = \{x \in [0,1]^2, |g(x)| \leq \tilde{\epsilon}\}$. Then*

$$\int_{X_{\tilde{\epsilon}}} |\nabla g(x)|\, dx \leq C \cdot \tilde{\epsilon},$$

where $C = 2 \sup_{t \in \mathbb{R}} H^1(\partial X_t)$ and H^1 is the Hausdorff measure, or length.

Proof. See the Appendix. □

We have shown that $N_B \leq C \cdot 2^j$. Let us now translate this result into a separation rank for the kernel $K(\xi, \eta, t)$, by means of the smooth partition of unity $q_k(x)$ already alluded to earlier in this section. Specifically, take a C^∞ function $\chi(x)$ such that $\chi(x) > 0$ for $|x| < 1$ and $\chi(x) = 0$ for $|x| \geq 1$. Consider the collection x_k of all ball centers, including those outside the set $X_\mu^{\mu'}(\delta)$. Then for each x_k define

$$\tilde{q}_k(x) = \chi\left(\frac{x - x_k}{\rho_k}\right).$$

By Lemma 4.2, each $\tilde{q}_k(x)$ has ϵ-separation rank of order $O(\epsilon^{-1/M})$ for all $M > 0$. The partition of unity is then, in the usual manner, defined as

$$q_k(x) = \frac{\tilde{q}_k(x)}{\sum_k \tilde{q}_k(x)}.$$

The constant overlap property, valid in a neighborhood of $X_\mu^{\mu'}(\delta)$, ensures that the smoothness constants of $q_k(x)$ are comparable to those of $\tilde{q}_k(x)$, as long as x_k is in or near $X_\mu^{\mu'}(\delta)$. As a matter of illustration, Lemma 4.2 would apply to those $q_k(x)$ near $X_\mu^{\mu'}(\delta)$ and yields an ϵ-separation rank of order $O(\epsilon^{-1/M})$ for all $M > 0$. (In truth, we will apply Lemma 4.2 later to a more complicated amplitude involving $q_k(x)$.)

At this point, recall that we are trying to separate the restricted kernel (4.30) on supp $\chi_\mu \times$ supp $\chi_{\mu'}$, that we have linearized the phase in η and that we are linearizing it in x as in equation (4.34). The point of $q_k(x)$ is that the quadratic contribution can be absorbed in the amplitude without changing the symbol properties of the latter (equation (4.33)). The new amplitude a_k is defined from

$$q_k(x) a_k(x, \eta, t) \chi_{\mu'}(\eta) = q_k(x) e^{\frac{i}{2}(x - x_k)^t \nabla_x \nabla_x \eta \cdot \nabla_\eta \Phi(y(x), \eta_0, t)(x - x_k)} a(x, \eta, t) \chi_{\mu'}(\eta).$$

The constant and linear contributions to the phase are

$$\eta \cdot \nabla_\eta \Phi(x_k, \eta_0, t) + (x - x_k) \cdot \nabla_x (\eta \cdot \nabla_\eta \Phi(x_k, \eta_0, t)) - x \cdot \xi.$$

The first term, $\eta \cdot \nabla_\eta \Phi(x_k, \eta_0, t)$, and the third term $-x_k \cdot \nabla_x (\eta \cdot \nabla_\eta \Phi(x_k, \eta_0, t))$ are both independent of x and separable in η, so we can ignore them. What remains is a modified kernel of the form

$$K_k(\xi, \eta, t) = \int q_k(x) e^{ix \cdot (A(t)\eta - \xi)} a_k(x, \eta, t)\, dx, \tag{4.41}$$

where $A(t) = \nabla_x \nabla_\eta \Phi(x_k, \eta_0, t)$. For sufficiently small times, that is $t = O(2^{-j/2})$, it turns out that K_k "looks enough like a pseudodifferential operator" and has constant ϵ-separation rank. When t gets larger than $2^{-j/2}$, this property quickly degrades, however. In order to justify these claims, consider the changes of variables

$$x' = \frac{x - x_k}{\rho_k}, \qquad \xi' = \frac{\xi - \xi_\mu}{2^j}, \qquad \eta' = \frac{\eta - \eta_{\mu'}}{2^j}. \tag{4.42}$$

Translations and dilations do not affect separation ranks. Their effect is to normalize the kernel so that the integral in x' is in a region of size at most $O(1)$ in x, and the range for ξ' and η' is a ball centered at the origin, with $O(1)$ radius. The new amplitude

$$b_k(x', \eta', t) = q_k(x(x')) a_k(x(x'), \eta(\eta'), t) \chi_{\mu'}(\eta(\eta'))$$

is a C^∞ function whose smoothness constants do not depend on j or j' anymore, because in the new variables, the symbol conditions (4.33) read

$$|\partial_{\eta'}^\alpha \partial_{x'}^\beta b_k(x', \eta', t)| \le C_{\alpha,\beta} 2^{j|\alpha|} (1 + |\eta_{\mu'} + 2^j \eta'|)^{-|\alpha|/2} \le C_{\alpha,\beta}.$$

(We have used $|\eta_{\mu'}| \asymp 2^{2j}$.) As for the phase, we have $A(t) = I + tP(t)$ by Lemma 4.3, with $P(t) = O(1)$ componentwise. Therefore,

$$x \cdot (A(t)\eta - \xi) = \rho_k 2^j x' \cdot (\eta' - \xi') + t\rho_k 2^j x' \cdot (P(t)\eta' - \xi') + \text{OK}. \tag{4.43}$$

The term "OK" refers to quantities that depend either on x', or on (η', ξ') – but not on all three at the same time, hence absorbable in the amplitude.

Let us now distinguish three subcases, depending on how t asymptotically compares to $2^{-j/2}$. Recall that $C2^{-j} \le \rho_k \le C2^{-j/2}$.

4.4.2.1 Typical Times, $2^{-j} \lesssim t \lesssim 2^{-j/2}$

If $t \lesssim 2^{-j/2}$, then $t\rho_k 2^j \le C$ and hence the second term in (4.43) is non-oscillatory and can be absorbed in the amplitude b in a now standard manner. What remains is

$$\int e^{i\rho_k 2^j x' \cdot (\eta' - \xi')} b_k(x', \eta', t) dx' = (2\pi)^2 \hat{b}_k(\rho_k 2^j (\eta' - \xi'), \eta')$$

and can be seen to have ϵ-separation rank $O(\epsilon^{-1/M})$, by applying Lemma 4.2 to the properly supported C^∞ function $\hat{b}_k$ (the diagonal scaling by $\rho_k 2^j$ is harmless.) The overall separation rank is proportional to the number of balls used to cover the set $X_\mu^{\mu'}(\delta)$, hence of order $O(2^j)$, as claimed in Theorem 4.1.

Note that Lemma 4.2 should actually be invoked with an adequate fraction of ϵ, to make sure that

$$\|R_2\|_{L^2 \to L^2} \le \frac{\epsilon}{4}. \tag{4.44}$$

In the appendix we settle an inconspicuous complication arising in the justification of (4.44), having to do with the fact that the separation remainder is actually a sum over $O(2^j)$ contributions, as in equation (4.31).

An application of Lemma 4.4 now shows that each wave atom submatrix formed from R_2, with indices $\mathbf{n}$, $\mathbf{n}'$, has $\ell^2_{\mathbf{n}'} \to \ell^2_{\mathbf{n}}$ norm at most $\epsilon/4$.

4.4.2.2 Large Times, $t \gtrsim 2^{-j/2}$

If asymptotically $t \ge 2^{-j/2}$ then $t\rho_k 2^j$ grows in j and a different definition of $q_k(x)$ is necessary. More precisely, we repeat the covering argument of $X_\mu^{\mu'}(\delta)$ with a smaller local ball radius density, given by

$$\rho(x, j) = \frac{1}{\sqrt{t^2 2^{2j} + 2^{2j} |\nabla g(x)|}}.$$

All ball radii now obey $\rho_k \le \frac{1}{t2^j}$, hence the phase becomes $\rho_k 2^j x' \cdot (\eta' - \xi')$ + non-oscillatory, as required. By repeating the previous counting argument, their total number is $O(t^2 2^{2j})$. The rest of the argument is otherwise identical.

The conclusion is the same as before: the overall separation rank is proportional to the number of balls used in the main partitioning argument, here $O(t^2 2^{2j})$. The justification that R_2 gives rise to submatrices of norm $\epsilon/4$ is the same as before.

4.4.2.3 Small Times, $\mathbf{t \lesssim 2^{-j}}$

For small times, the same argument would apply, but a major simplification of the problem's geometry allows us to prove a stronger result. By Lemma 4.3,

$$\nabla_x \Phi(x, \eta, t) = \eta + O(t|\eta|) = \eta + O(t 2^{2j}).$$

For $t \leq C \cdot 2^{-j}$ there exists a value of δ for which the set $X_\eta^\eta(\delta)$ defined by the condition $|\nabla_x \Phi(x, \eta, t) - \eta| \leq 2^j \delta$ covers $[0, 1]^2$. So will $X_\mu^{\mu'}(\delta)$, which is bigger than $X_\eta^\eta(\delta)$. The neighborhood of the locus of stationary phase is, therefore, the whole unit square.

We follow the same reasoning as before, and try to find a covering of $[0, 1]^2$ with balls of radius ρ_k in which the second-order term in the x-expansion of the phase is non-oscillatory. For $t = O(2^{-j})$ it suffices to take $\rho_k = \rho_0$, identically equal to

$$\rho_0 = \frac{1}{\sqrt{t} 2^j}.$$

Indeed, by Lemma 4.3,

$$\rho_0^2 |\nabla_x \nabla_x \Phi(t, \eta, t)| \leq C \cdot \rho_0^2 t |\eta| \leq C.$$

The collections of ball centers x_k can be taken as the Cartesian grid

$$x_k = (k_1, k_2) \frac{1}{2} t^{-1/2} 2^{-j}, \qquad k_1, k_2 \in \mathbb{Z}.$$

This corresponds to $O(1 + t 2^{2j})$ balls $B_{x_k}(\rho_0)$. The exact same reasoning as in the more general case applies, and yields an overall $\epsilon/4$-separation rank of order $O(1 + t 2^{2j})$.

4.4.3 Fine Scales, Nonstationary Phase

Let us now show that the nonstationary phase part yields a negligible contribution. Recall that we have defined, for each $\delta > 0$,

$$X_\mu^{\mu'}(\delta) = \{x \in [0,1]^2; \text{there exist } \xi \in \text{supp}\chi_\mu \text{and } \eta \in \text{supp}\chi_{\mu'},\ |\xi - \nabla_x \Phi(x,\eta,t)| \le 2^j \delta\}. \tag{4.45}$$

The partition of unity $\{q_k(x)\}$ introduced in the previous section can be used as smooth indicators for the complement of $X_\mu^{\mu'}(\delta)$. Let S_{out} be the set $\{x_k : B_{x_k}(\rho_k) \cap X_\mu^{\mu'}(\delta) \neq \emptyset\}$, and

$$p(x) = \sum_{x_k \in S_{\text{out}}} q_k(x).$$

Of course, $p(x)$ depends on j, j' and δ but keeping track of this fact would make the notations unnecessarily heavy. It follows from the definition of $q_k(x)$ that we have the "maximal" smoothness condition

$$\sup_{x \in [0,1]^2} |\partial_x^\alpha p(x)| \le C_\alpha \cdot 2^{j|\alpha|}.$$

We can now readily estimate

$$R_3(\xi,\eta,t) = \int e^{i(\Phi(x,\eta,t) - x\cdot\xi)}(1 - p(x))a(x,\eta,t)\,dx.$$

Indeed, we claim that an adequate choice of δ implies $\|R_3\|_2 \le \epsilon/4$ in L^2. To this end, let us first check L^2 boundedness. The smoothness property of $p(x)$, along with the estimate $2^{2j} \sim |\xi|$, imply that the amplitude

$$\sigma(x,\eta,t) = (1 - p(x))a(x,\eta,t)$$

is a symbol of order zero and type $(1, 1/2)$, in the sense that

$$|\partial_\xi^\alpha \partial_x^\beta \sigma(x,\xi,t)| \le C_{\alpha,\beta}(1 + |\xi|)^{-|\alpha| + |\beta|/2}.$$

As mentioned earlier, it is a beautiful application of the wave atom sparsity Theorem that Fourier integral operators of type $(1/2, 1/2)$, and in particular the kernel $e^{i\Phi}\sigma$ with σ as

above, are bounded on L^2.

Let us now show that the L^2 bound can be made arbitrarily small, by an adequate choice of δ. Consider the differential operator

$$L = \frac{1}{|\xi - \nabla_x \Phi(x,\eta,t)|^2} \, (\Delta_x - i\Delta_x \Phi(x,\eta,t) I),$$

which is chosen so that $L e^{i(\Phi(x,\eta,t) - x\cdot\xi)} = e^{i(\Phi(x,\eta,t) - x\cdot\xi)}$. The operator L can be applied any number of times to the exponential factor, and then moved to $\sigma = (1-p)a$ by integration by parts. The effect on the amplitude σ is the following:

- Every $\frac{1}{|\xi - \nabla_x \Phi(x,\eta,t)|^2}$, on the support of $(1-p(x))\chi_{\mu'}(\eta)$, brings in a factor $\frac{1}{\delta^2 2^{2j}}$, thanks to the definition of the set $X_\mu^{\mu'}(\delta)$.
- Every $L(1-p(x))$ yields a factor 2^{2j}, because of the smoothness property of $p(x)$.
- Every $\Delta_x \Phi$ yields a factor 2^{2j}, by homogeneity.
- After integration by parts, the new amplitude obeys the same smoothness assumptions as σ, hence is still a symbol of type $(1, 1/2)$.

Therefore, we conclude that

$$\delta^{2M} L^M \sigma(x,\eta,t)$$

is of type $(1, 1/2)$, with smoothness constants depending on M, but independent of δ. Invoking the general theory of FIOs, the L^2 bound on R_3 is therefore of the form

$$\|R_3\|_2 \leq C_M \delta^{-2M} \tag{4.46}$$

For fixed M, this bound can be made less than $\frac{\epsilon}{8\pi}$ by choosing δ as

$$\delta \geq C'_{M'} \epsilon^{-\frac{1}{M'}}, \tag{4.47}$$

with $M' = 2M$ and for some constant $C'_{M'}$ related to C_M. The combination of this result and Lemma 4.4 translates into a boundedness result for the corresponding submatrix in $\mathbf{n}$, $\mathbf{n}'$, namely that its ℓ^2 norm is bounded by $\epsilon/4$.

The proof is now complete, because the remainders R_1, R_2 and R_3 are of size at most $\epsilon/2$, $\epsilon/4$ and $\epsilon/4$ respectively, hence add up to ϵ. $\square$

4.5 Scattering Estimates

The objective of this section is to quantify the interactions, or energy transfer from an input wave vector $\xi_{\mu'}$ to other output wave vectors ξ_μ. As a result, we will obtain estimates on the *sum* of ranks of submatrices, either on $j, \mathbf{m}, \nu$ or $j', \mathbf{m}', \nu'$.

Theorem 4.2. *Let $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$ be the submatrix corresponding to $j, \mathbf{m}, \nu$ and $j', \mathbf{m}', \nu'$ in the separated wave atom representation of $E(t)$. For any $\epsilon > 0$, given $(j', \mathbf{m}')$, let $\Omega_{j',\mathbf{m}'}(t)$ be the smallest set of wave vectors $(j, \mathbf{m})$ such that setting $E(t)_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'} = 0$ for $(j, \mathbf{m}) \notin \Omega_{j',\mathbf{m}'}(t)$ and all ν, ν' results in an error less than ϵ in matrix ℓ_2 norm. Then the cardinality of $\Omega_{j',\mathbf{m}'}(t)$ obeys the bound*

$$|\Omega_{j',\mathbf{m}'}(t)| \leq C_\epsilon \cdot (1 + t^2 2^{2j'}),$$

where $C_\epsilon \leq C_M \epsilon^{-1/M}$, for all $M > 0$.

Proof. Fix $\epsilon > 0$ and $M > 0$. For this proof, we will exploit the compression properties of the wave propagator as in Theorem 1.1. The wave atom representation $\tilde{E}_{B,N}(t)$ of the propagator $E(t)$ is constructed as a matrix with two shifted band diagonals indexed by $\nu = \pm$, each of them corresponding to a ball in phase space centered about $h_{t,\nu}(\mu')$, and defined through the wave atom metric ω. More precisely, the "shifted band diagonals" are defined as the following set of wave atom subscripts:

$$SBD(\mu') = \bigcup_{\nu=\pm} \{\mu : \omega(\mu, h_{t,\nu}(\mu')) \leq r\},$$

with r chosen such that $|\{\mu : \omega(\mu, h_{t,\nu}(\mu')) \leq r\}| \asymp B$. Take B large enough so that the right-hand side of the error estimate (1.20) obeys $C_M B^{-M} \leq \epsilon$. Then of course $r \leq C_M \epsilon^{-1/M}$. Note that an error ϵ in L^2 for operators translates into an error ϵ in ℓ^2 for the wave atom matrix, by the tight frame property. In turn, restriction to a certain subset of rows and columns implies an error smaller than ϵ in ℓ^2 for the submatrix corresponding to $j, \mathbf{m}, \nu$ and $j', \mathbf{m}', \nu'$.

To estimate the size of $\Omega_{j',\mathbf{m}'}(t)$ it suffices to count the number of wave vectors $(j, \mathbf{m})$ which are part of at least one element $\mu = (j, \mathbf{m}, \mathbf{n})$ of the union of the shifted band diagonals

over all $\mathbf{n}'$,

$$\bigcup_{\mathbf{n}':\mu'=(j',\mathbf{m}',\mathbf{n}')} SBD(\mu').$$

To this effect, recall that the local wave vector $\xi(t) = \nabla_x \Phi(t,x)$ is obtained from the solution of the Hamilton-Jacobi equation $\Phi_t = c(x)|\nabla_x \Phi|$ with initial condition $\Phi(0,x) = x \cdot \xi_{\mu'}$. The range of all such wave vectors defines a region in the frequency plane, which can be inscribed in a ball Q_0 centered at $\xi_{\mu'}$ and of radius majorized by $C \cdot t|\xi_{\mu'}| \leq C \cdot t 2^{2j'}$. The set of wave vectors ξ_μ defined through $SBD(\mu')$ is slightly larger however, because the radius r is nonzero, but can certainly be inscribed in a larger ball Q_r of radius bounded by $C \cdot t 2^{2j'} + C_\epsilon 2^{j'}$.

It remains to count the number of tiling indicators $\chi_\mu(\xi)$ whose supports intersect the ball Q_r. Near $\xi_{\mu'}$, the support of each indicator has radius $O(2^{j'})$, so it suffices to use a number of indicators bounded by

$$C \cdot \left(\frac{C \cdot t 2^{2j'} + C_\epsilon \cdot 2^{j'}}{2^{j'}} \right)^2 \leq C_\epsilon \cdot (1 + t^2 2^{2j'}).$$

This is the desired bound on the cardinality of $\Omega_{j',\mathbf{m}'}(t)$. □

A simple counting argument now allows us to formulate the following result, companion to Theorem 4.1. The collection of bounds is summarized in Figure 4.11.

Corollary 4.1. *Consider the submatrix $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$ obtained by fixing $(j,\mathbf{m})$ and $(j',\mathbf{m}')$ in the wave atom representation of the propagator $E(t)$ after reordering $(n_1,n_2;n_1',n_2') \to (n_1,n_1';n_2,n_2')$. Denote by $r_{j\mathbf{m}}^{j'\mathbf{m}'}$ the maximum over ν,ν' of the ϵ-rank of $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$. Then we have the bounds*

- *for $t \lesssim 2^{-j}$, $\sum_{j\mathbf{m}} r_{j\mathbf{m}}^{j'\mathbf{m}'} \leq C_\epsilon \cdot (1 + t 2^{2j})$,*
- *for $2^{-j} \lesssim t \lesssim 2^{-j/2}$, $\sum_{j\mathbf{m}} r_{j\mathbf{m}}^{j'\mathbf{m}'} \leq C_\epsilon \cdot t^2 2^{3j}$,*
- *for $2^{-j/2} \lesssim t \leq T$, $\sum_{j\mathbf{m}} r_{j\mathbf{m}}^{j'\mathbf{m}'} \leq C_\epsilon \cdot 2^{2j}$,*

with $C_\epsilon \leq C_M \epsilon^{-1/M}$, for all $M > 0$, and C_ϵ also depends on T. The same bounds are valid for $\sum_{j'\mathbf{m}'} r_{j\mathbf{m}}^{j'\mathbf{m}'}$.

Proof. For $t \leq 2^{-j/2}$, or a constant multiple thereof, we can combine Theorem 4.1 with the scattering estimate (4.2) to obtain the first two bounds. For $t \geq 2^{-j/2}$, it suffices to notice

that the rank of each submatrix $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$ must be smaller than the number of nonzero elements. After thresholding at level ϵ in ℓ^2, the number of nonzero elements in any of the matrices $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$, for fixed $\mathbf{m}'$, is bounded by $C_\epsilon \cdot 2^{2j}$, by sparsity (Theorem 1.1). The third bound follows.

The same bounds on $\sum_{j'\mathbf{m}'} r_{j\mathbf{m}}^{j'\mathbf{m}'}$ stem from the observation that the adjoint operator $E^*(t)$ is obtained from the backward-in-time wave equation, which admits the same sparsity and separation properties. Note that formulating bounds in terms of j or j' does not make any difference since $j \asymp j'$ by sparsity.

□

4.6 Special Cases

In this section we continue the study of three of the four representative sample media introduced in Section 4.3, as well as another medium called "misaligned wave guide," this time in the light of the rank estimates just obtained. In two cases (Wave Guide and Bumps) the rank and complexity estimates turn out to be quite pessimistic and we are able to prove better bounds under certain conditions. In the two other less favorable cases (Misaligned Wave Guide and Linear Mirror), we give heuristic arguments that the rank bounds of Section 4.4 and 4.5 are in fact attained.

4.6.1 Wave Guides

We refer to a wave guide as an acoustic medium whose speed of sound depends only on one coordinate, either x_1 or x_2. As always, it is also assumed to be C^∞.

The rank bounds can be significantly improved for wave guides. In short, we show that rank majorants for wave guides are in general the *square root* of the rank majorants in the worst case.

Theorem 4.3. *Assume the velocity profile depends only on x_2 and is C^∞. Consider the submatrix $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$ obtained by fixing $(j,\mathbf{m})$ and $(j',\mathbf{m}')$ in the wave atom representation of the propagator $E(t)$, after reordering $(n_1,n_2;n_1',n_2') \to (n_1,n_1';n_2,n_2')$. Then the ϵ-rank of $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$ obeys*

- *for $t \lesssim 2^{-j}$, $r \le C_\epsilon \cdot (1+\sqrt{t}2^j)$,*

- *for* $2^{-j} \lesssim t \lesssim 2^{-j/2}$, $r \leq C_\epsilon \cdot 2^{j/2}$,
- *for* $2^{-j/2} \lesssim t \leq T$, $r \leq C_\epsilon \cdot t2^j$,

with $C_\epsilon \leq C_M \epsilon^{-1/M}$, *for all* $M > 0$, *and* C_ϵ *also depends on* T.

Proof. When the velocity profile $c(x)$ does not depend on x_1, it is easy to check that the local wave numbers $\nabla_x \Phi_\pm(x, \xi, t)$ do not depend on x_1 either (although $\Phi_\pm$ itself does). The steps of the proof are then the same as for Theorem 4.1, except that the definition of indicators $q_k(x)$ is a bit different. Instead of considering balls $B_{x_k}(\rho_k)$, we will consider horizontal *strips* $S_{x_{2,k}}(\rho_k)$ centered at height $x_2 = x_{2,k}$ and of width $2\rho_k$. Equations (4.34) through (4.37) then carry through unchanged, but a major simplification occurs in the counting argument for N_S, the number of strips necessary to make the restrictions of the phase non-oscillatory on each $q_k(x)$. The problem is now one-dimensional, g depends on x_2 only, so the local "strip width density" can be defined as

$$\rho(x_2, j) = \frac{1}{\sqrt{2^j + 2^{2j}|g'(x_2)|}}, \tag{4.48}$$

and the lattice y_k can be replaced by a simpler one-dimensional sequence $y_{2,k} = kb2^{-j}$. In contrast with equation (4.40), the weights w_k assigned to $y_{2,k}$ must now be defined as

$$w_k = 2^{-j} \sum_{x_{2,\ell} \in X : y_{2,k} \in S_{x_{2,\ell}}(\rho_\ell)} \frac{1}{\rho_\ell}.$$

There are $O(2^j \rho_\ell)$ points $y_{2,k}$ inside the interval $[x_{2,\ell} - \rho_\ell, x_{2,\ell} + \rho_\ell]$, so we have

$$N_S \leq C \cdot \sum_k w_k.$$

The corresponding integral is

$$N_S \leq C \cdot \int_X \frac{1}{\rho(x_2, j)}\, dx_2.$$

We should now use (4.48) in combination with the bound $\sqrt{2^j + 2^{2j}|g'(x_2)|} \leq 2^{j/2}(1 + \frac{1}{2}2^j|g'(x_2)|)$ and Lemma 4.6—also valid in dimension one—to obtain the improved bound

$$N_S \leq C \cdot 2^{j/2}.$$

The rest of the proof then proceeds in an analogous way.

- For typical times $2^{-j} \lesssim t \lesssim 2^{-j/2}$, the bound on r is the same as that for N_S, namely $O(2^{j/2})$.
- For small times $t \lesssim 2^{-j}$, the strip heights can be taken equispaced and equal to

$$x_{2,k} = kbt^{-1/2}2^{-j},$$

  yielding $N_S = O(1 + \sqrt{t}2^j)$ strips and a comparable rank r.
- For large times $t \gtrsim 2^{-j/2}$,

$$\rho(x_2, j) = \frac{1}{\sqrt{t^2 2^{2j} + 2^{2j}|g'(x_2)|}},$$

  so $r \simeq N_S = O(t2^j)$ by the previous argument.

□

These rank bounds are summarized in Figure 4.11. Let us remark at this point that the rank plateaus at a value $O(2^j)$ for $t \simeq 1$, although the size of the matrix is $\simeq 2^{2j}$-by-2^{2j}. This is obviously a consequence of the above theorem for times before caustics, but it turns out the same result is also valid *after* caustics start forming. The justification of this more general claim will follow from the analysis of the stronger bound on the sum of ranks over j' and $\mathbf{m}'$, which we now present.

Theorem 4.4. *Assume the velocity profile depends only on x_2 and is C^∞. Consider the submatrix $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$ obtained by fixing $(j, \mathbf{m})$ and $(j', \mathbf{m}')$ in the wave atom representation of the propagator $E(t)$, after reordering $(n_1, n_2; n_1', n_2') \to (n_1, n_1'; n_2, n_2')$. Denote by $r_{j\mathbf{m}}^{j'\mathbf{m}'}$ the maximum over ν, ν' of the ϵ-rank of $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$. Then we have the bounds*

- *for $t \lesssim 2^{-j}$, $\sum_{j\mathbf{m}} r_{j\mathbf{m}}^{j'\mathbf{m}'} \le C_\epsilon \cdot (1 + \sqrt{t}2^j)$,*
- *for $2^{-j} \lesssim t \lesssim 2^{-j/2}$, $\sum_{j\mathbf{m}} r_{j\mathbf{m}}^{j'\mathbf{m}'} \le C_\epsilon \cdot t2^{3j/2}$,*
- *for $t \gtrsim 2^{-j/2} \le T$, $\sum_{j\mathbf{m}} r_{j\mathbf{m}}^{j'\mathbf{m}'} \le C_\epsilon \cdot 2^j$,*

with $C_\epsilon \le C_M \epsilon^{-1/M}$, for all $M > 0$, and C_ϵ also depends on T. The same bounds are valid for $\sum_{j'\mathbf{m}'} r_{j\mathbf{m}}^{j'\mathbf{m}'}$.

Proof. We need a stronger version of the scattering estimate in Theorem 4.2, in the special case of wave guides. The question is to determine the number of balls of radius $\simeq 2^j$ (each containing a wave atom bump in frequency) necessary to cover the locus of local wave vectors $\nabla_x \Phi_\pm(x,\xi,t)$, when a union is taken over all possible values of x. We know from the general case that this local wave vector cannot wander too far off ξ, namely $|\xi - \nabla_x \Phi_\pm(x,\xi,t)| \le C \cdot t 2^{2j}$, resulting in a covering by at most $O(1 + t^2 2^{2j})$ balls.

In the case of wave guides, however, this locus is for each phase a *one-dimensional smooth curve* Γ_ξ, generated by the union of all wave vectors over the single coordinate x_2 (because the local wave vector is independent of x_1). In addition, Γ_ξ inherits the homogeneity of degree one of Φ, which makes it homothetic in $|\xi|$. As a result, the length of Γ_ξ is in fact comparable to the diameter of the locus in the general case, $O(t2^{2j})$, so it only takes $O(t2^j)$ balls of radius $\simeq 2^j$ to cover Γ_ξ. As a result, the cardinality of the set of participating wave vectors, in analogy with Theorem 4.2, is

$$|\Omega_{j',\mathbf{m}'}(t)| \le C_\epsilon \cdot (1 + t2^j).$$

The argument bounding sums of ranks over j' and $\mathbf{m}'$ then goes on to follow from the proof of Corollary 4.1, and we obtain

- for $t \lesssim 2^{-j}$, $\sum_{j\mathbf{m}} r_{j\mathbf{m}}^{j'\mathbf{m}'} \le C_\epsilon \cdot (1 + \sqrt{t}2^j)$,
- for $2^{-j} \lesssim t \lesssim T$, before caustics, $\sum_{j\mathbf{m}} r_{j\mathbf{m}}^{j'\mathbf{m}'} \le C_\epsilon \cdot t 2^{3j/2}$.

Since we have so far relied on the existence of the phase functions $\Phi_\pm$ in our reasoning, we took the precaution of mentioning that the result is valid before the formation of caustics (on plane wave initial conditions). The same bounds also hold when the sum is taken over $(j', \mathbf{m}')$ instead of $(j, \mathbf{m})$, for the same reasons as previously.

We however claim that a stronger estimate holds: $\sum_{j',\mathbf{m}'} r_{j,\mathbf{m}}^{j',\mathbf{m}'} \le C_\epsilon \cdot 2^j$, regardless of t, even after caustics. This improves on the earlier bound when $t \gtrsim 2^{-j/2}$. In order to justify this claim, we need to understand the effect of the wave guide structure on the submatrices of interest, $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$ with row index $\mathbf{n} = (n_1, n_2)$ and column index $\mathbf{n}' = (n_1', n_2')$. For short, when the other parameters are encumbering, we also denote the submatrix by $E_{n_1,n_2}^{n_1';n_2'}$.

The subscript ν takes on two values ($\pm$) so we omit it in what follows. Recall the central sparsity result, Theorem 1.1, which states that for fixed $\mu = (j, \mathbf{m}, \mathbf{n})$ the number of

matrix elements above a threshold ϵ (in absolute value), spanned by the remaining indices $(j', \mathbf{m}', \mathbf{n}')$, is a constant $C_\epsilon = O(\epsilon^{-1/M})$ for all $M > 0$. Let us now make the exercise of only fixing $(j, \mathbf{m}, n_1)$: the number of elements above ϵ spanned by the other indices $(j', \mathbf{m}', \mathbf{n}', n_2)$ is proportional to the number of n_2's, that is $C_\epsilon \cdot 2^j$. Fixing n_1 means considering only a subset of the rows, i.e. "mutilating" each submatrix $E^{n_1', n_2'}_{n_1, n_2}$. Surely, for fixed n_1 the sum of ranks of those mutilated matrices over $j', \mathbf{m}'$ cannot exceed the total number of elements, $C_\epsilon \cdot 2^j$. Re-ordering the submatrices as $E^{n_2, n_2'}_{n_1, n_1'}$ does not change that fact.

As we now consider different values of n_1 (still for fixed $j, \mathbf{m}$), we introduce *no new information*. Because of the invariance of the problem under translations in x_1, we obtain the same wave atom matrix elements, albeit shifted circularly in n_1'. More precisely, the invariance property reads

$$E^{n_1', n_2'}_{n_1, n_2} = E^{n_1'+p, n_2'}_{n_1+p, n_2},$$

where p is any integer and addition is understood modulo the bound on the number of n_1. Consequently, the rank of $E^{n_1', n_2'}_{n_1, n_2}$ does depend on whether it is mutilated to a certain subset of n_1's or not. The same is true for the sum of ranks over $(j', \mathbf{m}')$, so the claim follows.

Again, the same bounds also hold when the sum is taken over $(j, \mathbf{m})$ instead of $(j', \mathbf{m}')$. The proof is complete.

□

4.6.2 Bumps

The example "Bumps" belongs to a larger class of nondegenerate oscillating profiles, which can be formalized as follows.

Definition 4.3. *(Transversality) A smooth velocity profile $c(x) > 0$ is said to be* transversal *when the following two conditions are satisfied:*

1. *The locus where the Hessian $\nabla\nabla c$ is singular is the union of a finite number of smooth curves.*

2. *For every point x for which there exists two unit vectors d, d' such that $(d\cdot\nabla)^2 c(x) = 0$ and $(d' \cdot \nabla)^3 c(x) = 0$, we have $d \cdot d' \neq 0$.*

As can easily be checked, examples of transversal profiles include smooth and separable functions $c(x_1, x_2) = \gamma_1(x_1)\gamma_2(x_2) > 0$ with γ_k''' nonzero when γ_k'' vanishes, $k = 1$ or 2. In the

"Bump" example, we have taken $\gamma_1(x) = \gamma_2(x) = \frac{3+\sin(4\pi x)}{16}$. We also expect a sum of wide bumps with random location and random positive amplitude to satisfy the transversality condition with high probability.

A notable example of non-transversal profile, on the other hand, would be the innocent-looking

$$c(x_1, x_2) = 2 + \sin(2\pi x_1)\sin(2\pi x_2),$$

for which condition 2 in definition 4.3 is violated.

The rationale for introducing "transversal" profiles is the following (obvious) asymptotic relation for the phase Hessian,

$$\nabla_x \nabla_x \Phi_\pm(x, \xi, t) = \pm t \nabla\nabla c(x)\,|\xi| + O(t^2|\xi|).$$

For small times $t = o(1)$, the locus of singularity of $\nabla_x \nabla_x \Phi_\pm$ is a deformation of that of $\nabla\nabla c$. Such information allows to characterize the locus $X^\eta_\xi(\delta)$ of stationary phase in a much more precise way than was done in the proof of theorem 4.1. As a result, the rank estimates can be strengthened as follows. The results are reported in Figure 4.11.

Theorem 4.5. *Assume $c(x)$ is smooth and transversal, in the sense of definition 4.3. Consider the submatrix $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$ obtained by fixing $(j, \mathbf{m})$ and $(j', \mathbf{m}')$ in the wave atom representation of the propagator $E(t)$, after reordering $(n_1, n_2; n_1', n_2') \to (n_1, n_1'; n_2, n_2')$. Then the ϵ-rank of $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$ obeys*

- *for $t \lesssim 2^{-j}$, $r \le C_\epsilon \cdot (1 + t2^{2j})$,*
- *for $2^{-j} \lesssim t \lesssim 2^{-j/3}$, $r \le C_\epsilon \cdot \frac{2^{j/2}}{\sqrt{t}}$,*
- *for $2^{-j/3} \lesssim t \le T = o(1)$, $r \le C_\epsilon \cdot t2^j$,*

with $C_\epsilon \le C_M \epsilon^{-1/M}$, for all $M > 0$.

Proof. As alluded to earlier, the condition $T = o(1)$ ensures that the phases $\Phi_\pm$ satisfy the same transversality conditions as $c(x)$.

The proof of the rank bound for $t \lesssim 2^{-j}$ is the same as previously, so let us consider $t \gtrsim 2^{-j}$. As alluded to earlier, the condition $T = o(1)$ ensures that the phases $\Phi_\pm$ satify the same transversality conditions as $c(x)$. The purpose of the transversality condition is

to allow a much more explicit description of the loci $X_\mu^{\mu'}(\delta)$ of stationary phase than in the proof of Theorem 4.1.

Consider one phase function, say $\Phi = \Phi_+$. Given a wave number η and a point $x^* \in [0,1]^2$, only three scenarios can occur.

1. Assume $\nabla_x \nabla_x \Phi(x^*, \eta, t) = 0$. By transversality, we necesarily have

$$|(d \cdot \nabla)^3 \Phi(x^*, \xi, \eta)| \geq C_{\text{trans}} t|\eta|, \tag{4.49}$$

uniformly over all unit vectors d. Let $\xi_0 = \nabla_x \Phi(x^*, \eta, t)$. We would like to find good bounds for the set

$$X_{\xi_0}^{\eta}(\delta) = \{x \in [0,1]^2 : |\nabla_x \Phi(x, \eta, t) - \xi_0| \leq \delta 2^j\}.$$

Once this is done, we can identify the wave atom subscripts μ, μ' such that ξ_μ is closest to ξ_0, η_μ is closest to η_0, and assert that $X_\mu^{\mu'}(\delta)$ has about the same size, up to a constant, as $X_{\xi_0}^{\eta}(\delta)$. See the reasoning leading to equation (4.32).

Using a Taylor expansion around x^* and Lemma 4.3 we first obtain

$$\nabla_x \Phi(x, \eta, t) = \xi_0 + \frac{1}{2} \sum_{k_1, k_2} (x - x^*)_{k_1} (x - x^*)_{k_2} \frac{\partial^2}{\partial x_{k_1} \partial x_{k_2}} \nabla_x \Phi(x^*, \eta, t) + O(|x - x^*|^3 t|\eta|). \tag{4.50}$$

We can take the dot product of this relation with $d(x) = \frac{x - x^*}{|x - x^*|}$ to get

$$d(x) \cdot (\nabla_x \Phi(x, \eta, t) - \xi_0) = \frac{1}{2} |x - x^*|^2 (d(x) \cdot \nabla_x)^3 \Phi(x^*, \eta, t) + O(|x - x^*|^3 t|\eta|).$$

The magnitude of a gradient is certainly greater than the absolute value of any directional derivative, so

$$|\nabla_x \Phi(x, \eta, t) - \xi_0| \geq \frac{1}{2} C_{\text{trans}} |x - x^*|^2 t|\eta| - O(|x - x^*|^3 t|\eta|).$$

When $|x - x^*| = o(1)$ as the scale j or equivalently $|\eta| \simeq 2^{2j}$ grows, then the $O(|x - x^*|^3 t|\eta|)$ remainder is asymptotically negligible and the behavior of Φ near x^* is governed by its third spatial derivatives. If we let $x \in X_{\xi_0}^{\eta}(\delta)$ then the condition

defining the latter set implies

$$C \cdot t|\eta||x - x^*|^2 \leq 2^j \delta,$$

which in turn shows that $X^{\eta}_{\xi_0}(\delta)$ is included in a ball centered at x^*, with radius ρ_X proportional to $\frac{2^{-j/2}}{\sqrt{t}}$. As t asymptotically exceeds 2^{-j} we are indeed in the regime where $|x - x^*| = o(1)$, validating smallness of the Taylor remainder.

With this information on the extent of the set of near-stationary phase points, we are ready to repeat the ball counting argument of Section 4.4.2. The argument consists in exhibiting balls $B_{x_k}(\rho_k)$ over which the phase is non-oscillatory in the sense that for $x \in B_{x_k}(\rho_k)$, it holds that

$$\rho_k^2 |\nabla_x \nabla_x \Phi(x, \eta, t)| \leq C. \tag{4.51}$$

In the neighborhood of x^* the phase Hessian obeys, componentwise,

$$|\nabla_x \nabla_x \Phi(x, \eta, t)| \leq C \cdot t|\eta||x - x^*|,$$

which means that for $x \in X^{\eta}_{\xi_0}(\delta)$ we have

$$|\nabla_x \nabla_x \Phi(x, \eta, t)| \leq C \cdot 2^{3j/2} \sqrt{t}.$$

To satisfy the non-oscillation condition (4.51), it suffices to take the ball radii r_k uniformly equal to

$$\rho_k \simeq 2^{-3j/4} t^{-1/4}. \tag{4.52}$$

This choice corresponds to a covering of $X^{\mu'}_{\mu}(\delta)$ by N_B balls, where

$$N_B \leq C \cdot \left(\frac{\rho_X}{\rho_k} \right)^2 = C \cdot 2^{j/2} t^{-1/2}. \tag{4.53}$$

This bound on N_B will be interpreted later as a rank estimate, because the zero Hessian scenario turns out to be the worst case (largest bound on N_B.). To this end, we now intend to review and compare the other two scenarios.

2. Assume now that there exists a direction d along which

$$d \cdot \nabla_x \Phi(x^*, \eta, t) = 0 \qquad \text{but} \qquad d^\perp \cdot \nabla_x \Phi(x^*, \eta, t) \neq 0.$$

In the direction d, we can repeat the argument of scenario 1 to conclude that the spatial extent of $X_\mu^{\mu'}(\delta)$ is of order $\rho_X(d) = 2^{-j/2} t^{-1/2}$. In the direction $d^\perp$, the situation is simpler because the Taylor expansion of $\nabla_x \Phi$ is the usual

$$\nabla_x \Phi(x, \eta, t) = \xi_0 + \sum_k (x - x^*)_k \frac{\partial}{\partial x_k} \nabla_x \Phi(x^*, \eta, t) + O(|x - x^*|^2 t|\eta|).$$

Repeating the sequence of steps leading up to (4.52), we obtain instead

$$\rho_X(d^\perp) \simeq 2^{-j} t^{-1}.$$

It is straightforward to check that the phase is always non-oscillatory in the sense of (4.51) over balls of radius $\rho_k = 2^{-j} t^{-1}$. We conclude that $X_\mu^{\mu'}(\delta)$ can be covered by N_B balls $B_{x_k}(\rho_k)$, with

$$N_B \leq C \cdot \frac{\rho_X(d) \rho_X(d^\perp)}{\rho_k^2} = C \cdot 2^{j/2} \sqrt{t}. \tag{4.54}$$

For times $t = O(1)$ this bound is always smaller than (4.53), obtained in scenario 1.

3. Finally, assume that the phase Hessian is nonsingular. By the same argument as above, the set $X_\mu^{\mu'}(\delta)$ can be inscribed in a ball of radius $\rho_X \simeq 2^{-j} t^{-1}$, over which the phase is non-oscillatory, resulting in

$$N_B \leq C, \tag{4.55}$$

independently of j. This latter bound is always smaller than (4.53) for times $t \gtrsim 2^{-j}$.

The conclusion of the above analysis is that the worst-case scenario arises when the Hessian vanishes, for which $N_B \leq C \cdot 2^{j/2} t^{-1/2}$. Before translating this bound into a rank estimate, we must make sure that the off-diagonal linear term in the phase (see equation (4.43)) is itself non-oscillatory. Recall that the normalizing change of variables (4.42) for x'

was chosen so that $x' = O(1)$ as long as $x \in B_{x_k}(\rho_k)$. In our case, we can choose it as

$$x' = \frac{x - x_k}{2^{-3j/4} t^{-1/4}},$$

resulting in

$$x \cdot (A(t)\eta - \xi) = (2^j t^{-1})^{1/4} x' \cdot (\eta' - \xi') + (2^j t^3)^{1/4} x' \cdot (P(t)\eta' - \xi'),$$

where $A(t) = I + tP(t)$ (compare with (4.43)). The term involving $P(t)$ is of order $O(1)$ as long as $t \lesssim 2^{-j/3}$, therefore allowing to view the bound on N_B as a rank estimate. That is the content of the second bullet in Theorem 4.5.

For times $t \gtrsim 2^{-j/3}$, we resort to the same reasoning as previously, namely modifying the change of variables as

$$x' = \frac{x - x_k}{(t2^j)^{-1}}.$$

This choice imposes a covering of $X_\mu^{\mu'}(\delta)$ by balls of radius $\rho_k = 2^{-j} t^{-1}$, resulting in

$$N_B \leq C \cdot \left(\frac{2^{j/2} t^{-1/2}}{2^{-j} t^{-1}} \right)^2 = C \cdot t 2^j.$$

The corresponding rank estimate follows (bullet 3 in Theorem 4.5.) This concludes the proof.

□

The corresponding result for sums of ranks is the following.

Theorem 4.6. *Assume the velocity profile is tranversal and C^∞. Consider the submatrix $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$ obtained by fixing $(j, \mathbf{m})$ and $(j', \mathbf{m}')$ in the wave atom representation of the propagator $E(t)$, after reordering $(n_1, n_2; n_1', n_2') \to (n_1, n_1'; n_2, n_2')$. Denote by $r_{j\mathbf{m}}^{j'\mathbf{m}'}$ the maximum over ν, ν' of the ϵ-rank of $E_{j\mathbf{m}\nu;j'\mathbf{m}'\nu'}(t)$. Then we have the bounds*

- *for $t \lesssim 2^{-j}$, $\sum_{j\mathbf{m}} r_{j\mathbf{m}}^{j'\mathbf{m}'} \leq C_\epsilon \cdot (1 + t2^{2j})$,*
- *for $2^{-j} \lesssim t \lesssim 2^{-3j/5}$, $\sum_{j\mathbf{m}} r_{j\mathbf{m}}^{j'\mathbf{m}'} \leq C_\epsilon \cdot \frac{2^{j/2}}{\sqrt{t}}$,*
- *for $2^{-3j/5} \lesssim t \leq T = o(1)$, $\sum_{j\mathbf{m}} r_{j\mathbf{m}}^{j'\mathbf{m}'} \leq C_\epsilon \cdot t^2 2^{2j}$,*

with $C_\epsilon \leq C_M \epsilon^{-1/M}$, for all $M > 0$. The same bounds are valid for $\sum_{j'\mathbf{m}'} r_{j\mathbf{m}}^{j'\mathbf{m}'}$.

Proof. The justification is a combination of the bounds of Theorem 4.5 with a scattering estimate, counting the number of wave vectors $\xi_{\mu'}$ involved in each scenario on the phase Hessian (see the proof of Theorem 4.5.) Fix a wave vector ξ_μ. The count is as follows:

1. We claim that the locus where the Hessian $\nabla\nabla c$ is identically zero contains at most a finite number of points, in the case of transversal velocity profiles. Assume by contradiction that it is not the case. By compactness there exists a sequence of points x_i in $[0,1]^2$ converging to some limit $x^* \in [0,1]^2$, such that $x_i \neq x^*$ and $\nabla\nabla c(x_i) = 0$. Necessarily, by continuity, $\nabla\nabla c(x^*) = 0$. Denote $d_i = \frac{x_i - x^*}{|x_i - x^*|}$. Since the unit circle is compact, there exists a subsequence d_{i_j} converging to some $d \in S^1$. It is then a simple matter to check to check that $(d \cdot \nabla)\nabla\nabla c = 0$, contradicting the transversality condition in Definition 4.3.

   The same property transfers to the phase Hessian for times $t = o(1)$. Each point x where $\nabla_x \nabla_x \Phi_\pm(x, \xi_\mu, t)$ vanishes identically corresponds to one wave vector, $\xi_0 = \nabla_x \Phi_\pm(x, \xi_\mu, t)$. As a consequence, there are at most a constant number of wave vectors $\xi_{\mu'}$ which belong in scenario 1, yielding a total combined rank

$$\sum_{(j,\mathbf{m}) \in I} r_{j,\mathbf{m}}^{j',\mathbf{m}'} \leq C \cdot \max\{2^{j/2} t^{-1/2}, 2^j t\}. \tag{4.56}$$

2. For scenario 2, we directly obtain from the transversality condition that the locus $\mathcal{L}$ where the phase Hessian is singular is a one-dimensional manifold. So is the locus Γ_{ξ_μ} of wave vectors $\xi = \nabla_x \Phi_\pm(x, \xi_\mu, t)$, where $x \in \mathcal{L}$. As in the proof of Theorem 4.4, the intersection of Γ_{ξ_μ} with the "scattering" ball $B_{\xi_\mu}(Ct2^{2j})$ can be covered by at most $O(t2^j)$ indicators $\chi_{\mu'}(\xi)$. As a result, the sum of ranks over $(j, \mathbf{m})$ for scenario 2 is

$$\sum_{(j,\mathbf{m}) \in II} r_{j,\mathbf{m}}^{j',\mathbf{m}'} \leq C \cdot 2^j t \cdot 2^{j/2} t^{1/2} = C \cdot 2^{3j/2} t^{3/2}. \tag{4.57}$$

3. Scenario 3 corresponds to all the wave vectors ξ_μ that are left out from scenarios 1 and 2. By Theorem 4.2, there are at most $O(t^2 2^{2j})$ of them. Each of those wave vectors corresponds to a submatrix with rank bounded by a $O(1)$ constant, so the total count is

$$\sum_{(j,\mathbf{m}) \in III} r_{j,\mathbf{m}}^{j',\mathbf{m}'} \leq C \cdot 2^{2j} t^2. \tag{4.58}$$

It now remains to add equations (4.56), (4.57) and (4.58). The two last bullets in Theorem 4.6 follow from the observation that (4.56) is asymptotically dominant when $t \lesssim 2^{-3j/5}$, but (4.58) dominates when $t \gtrsim 2^{-3j/5}$.

□

4.6.3 Misaligned Wave Guide

A "misaligned wave guide" is an essentially one-dimensional profile $c(x)$ whose redundant coordinate is not aligned with x_1 or x_2. One such example is

$$c(x_1, x_2) = 2 - \cos(2\sqrt{2}\pi(x_1 - x_2)),$$

which depends only on $x_1 - x_2$. We take the precaution to name those profiles essentially one-dimensional, because they should also be smooth and periodic on the torus, a requirement incompatible with being a wave guide in other directions than vertical, horizontal, or diagonal at 45 degrees as above.

The performance of our solver on "misaligned wave guide" is rather poor so we chose not to report it in Section 4.3.

We intent to justify, albeit not in a rigorous manner, that misaligned wave guides probably saturate the rank bound $r \lesssim 2^j$ of Theorem 4.1, when $t \simeq 2^{-j}$. We hope that this example may help illustrate a central piece of the argument behind Theorem 4.1.

Locally near the diagonal $x_1 = x_2$, we have $c(x_1, x_2) \simeq 1 + 4\pi^2(x_1 - x_2)^2$. The phases $\Phi_\pm$ therefore obey the small-time (and small $|x_1 - x_2|$) asymptotic relations

$$\Phi_\pm(x, \xi, t) \simeq x \cdot \xi \pm t(1 + 4\pi^2(x_1 - x_2)^2)|\xi|.$$

Let us now explain why the most expensive contribution in the phase, in terms of the resulting ranks, is the off-diagonal term proportional to $tx_1x_2|\xi|$. We had already alluded to this fact in Section 4.4.2. We remind the reader that $|\xi| \simeq 2^{2j}$, so we will simply consider the phase $2^j x_1 x_2$.

In view of the proof of Theorem 4.1, we would like to bound the cardinality of a covering of the locus $X_\xi^\eta(\delta)$ of near-stationary phase by balls inside which the phase satisfies the stronger requirement of being non-oscillatory, see (4.35). For any given $\xi = \eta$ and large δ it is easy to see that the locus $X_\xi^\eta(\delta)$ actually covers the whole unit square. The phase

Hessian is

$$\nabla_x \nabla_x \Phi(x_1, x_2) = 2^j \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

which implies a uniform ball radius $\rho_k \simeq 2^{-j/2}$. It takes $O(2^j)$ balls of radius ρ_k to cover the whole unit square, resulting in the announced bound $r \simeq 2^j$ for the rank.

4.6.4 Linear Mirror

A "linear mirror" is a profile $c(x)$ which is locally of the form $C + x \cdot \lambda$ for some vector λ. Of course $x \cdot \lambda$ is not compatible with smoothness and periodicity on the torus; see Section 4.3 for a good compromise.

Linear mirrors are representative of a class of profiles for which the rank bound of Theorem 4.1 is expected to be sharp. Again, we will not provide a rigorous proof but only give indications towards this claim.

In the region where $c(x) = C + x \cdot \lambda$, the phases can be solved for explicitly,

$$\Phi_\pm(x, \eta, t) = x \cdot \eta \pm Ct|\xi| \pm x \cdot \frac{\lambda}{|\lambda|}|\eta|(e^{t|\lambda|} - 1).$$

In analogy with equation (4.43), this expression can be linearized in η and rewritten as $x \cdot A(t)\eta +$ OK. In our case, the matrix elements of $A(t)$ are, for small time t, given as

$$A_{ij}(t) = \delta_{ij} + t\frac{\lambda_i}{|\lambda|}\frac{\eta_j}{|\eta|} + O(t^2).$$

In the notations of Section 4.4.2, we identify $P_{ij} = \frac{\lambda_i}{|\lambda|}\frac{\eta_j}{|\eta|}$. This is a prototypical non-diagonal matrix. This example leads us to believe that the linear part of the phase genuinely affects the rank estimates, and that we are not in presence of a proof artifact.

4.7 Discussion

So far we have assumed periodic boundary conditions for the wave equation inside the unit square $[0, 1]^2$, but simple modifications will allow the wave atom algorithm to work in slightly more general settings.

First, we can consider standard boundary conditions like Dirichlet ($u = 0$ on the boundary) or Neumann ($\frac{\partial u}{\partial n} = 0$) in the same domain $[0, 1]^2$. The two cases can be handled in a

straightforward manner by mirror extension of the computational domain to the periodized square $[0,2]^2$ with velocity

$$\tilde{c}(x_1,x_2) = \begin{cases} c(x_1,x_2) & \text{if} \quad 0 \le x_1, x_2 < 1, \\ c(2-x_1,x_2) & \text{if } 1 \le x_1 < 2,\ 0 \le x_2 < 1, \\ c(x_1,2-x_2) & \text{if } 0 \le x_1 < 1,\ 1 \le x_2 < 2, \\ c(2-x_1,2-x_2) & \text{if} \quad 1 \le x_1, x_2 < 2. \end{cases}$$

The wave equation can then be solved up to some time T for $\tilde{u}$ in the periodized extended square $[0,2]^2$, and $\tilde{u}$ mirror folded back onto $[0,1]^2$ using the rule

$$u(x_1,x_2) = \tilde{u}(x_1,x_2) - \tilde{u}(2-x_1,x_2) - \tilde{u}(x_1,2-x_2) + \tilde{u}(2-x_1,2-x_2)$$

if u is to satisfy Dirichlet boundary conditions, or

$$u(x_1,x_2) = \tilde{u}(x_1,x_2) + \tilde{u}(2-x_1,x_2) + \tilde{u}(x_1,2-x_2) + \tilde{u}(2-x_1,2-x_2)$$

if u is to satisfy Neumann boundary conditions. Some other choice of signs are possible and would lead, for example, to Dirichlet on two opposite sides and Neumann on the two other sides. For the wave atom algorithm to perform accurately on the extended domain, we need to ensure sure that $\tilde{c}(x_1,x_2)$ remains sufficiently smooth after mirror extension as above.

The increase in complexity resulting from the doubling of N, the number of grid points per dimension, may however be unacceptable in some applications. Readers interested in a more elegant treatment of boundary considerations, in the context of some other basis of bandlimited functions (prolate spheroidal wavefunctions,) should refer to the recent work of Beylkin and Sandberg, [9].

More generally, if the computational domain can be mapped onto the unit square by means of a smooth diffeomorphism, then it is only a matter of changing variables and re-using the same algorithm on the transformed equation. More complicated geometries or topologies would pose a significant challenge to wave-packet-type methods and their treatment would go far beyond the scope of this thesis.

Finally, wave atoms seem to be a promising tool for implementing absorbing bound-

ary conditions in the regime of high-frequency solutions. Assume for a moment that the wavefield $u(t,x)$ has frequency support obeying $|\xi| \geq \lambda$, and that the profile $c(x)$ is near constant near the edges of the unit square. Then the computational domain can be extended to include a surrounding buffer strip of width $O(\frac{1}{\sqrt{\lambda}})$ and constant sound speed, in which outgoing wave atoms can be safely removed from the solution by putting the corresponding matrix elements to zero. This should work provided the upscaled time step τ is of order $\tau = O(\frac{1}{\sqrt{N}})$.

Chapter 5

Conclusion

5.1 Achievements

The main contribution of this thesis is perhaps the single message that thinking in terms of geometric compression definitely creates opportunities for exciting new developments in numerical analysis.

Our achievements are only a small part of that program. We showed that the Green's function of the wave equation in smooth media is represented as a sparse matrix in the *curvelet* frame, as well as in the new *wave atom* frame. Those are essentially the only two universal change of bases in which sparsity occurs.

Those new mathematical insights translate into efficient algorithms for the wave equation. We studied in detail an embodiment of the repeated squaring for the Green's function in which high-dimensional separation techniques in the wave atom domain play an essential role, complementary to sparsity. The new algorithm has spectral accuracy and sometimes competes favorably against a pseudo-spectral method when a given wave equation needs to be solved several times with different initial conditions. In the process, we developed fast discrete transforms of independent interest, for curvelet and wave atoms.

We call "time upscaling" the possibility of compressing the Green's function for times large than the CFL timestep, for the purpose of speeding up computations.

5.2 Outlook

The main Theorem 1.1 can be generalized in a variety of ways. The same sparsity question can be posed in regimes of reflection and refraction through smooth interfaces – discontinu-

ities of $c(x)$—and the answer is probably positive when the wavefield is 'microlocally away' from the interface, in a sense to be made precise.

We would regard as mathematically significant any result of conservation of curvelet sparsity for nonlinear wave equations, e.g., with a nonlinearity of the type u^3.

As for the curvelet transform, our architecture can be made more useful or attractive in a number of ways and we discuss two opportunities.

- First, the redundancy of our transform is about 2.8 when wavelets are chosen at the finest scale, and 7.2 otherwise. For certain image processing tasks, redundant transformations may be of benefit, but for others, digital transforms with low redundancy might be more desirable. It is not immediate how one could adapt our ideas to reduce the redundancy while keeping the isometry property and remaining faithful to the continuous transform. In particular, it is not known whether one can construct orthonormal bases of curvelets. We regard this problem as very significant and extremely challenging.

- Second, compactly supported (or at least exponentially decaying) curvelets would have the potential to yield sparser expansions of images with geometrical regularity. We consider the design of compactly supported curvelet tight frames as another interesting open problem.

Additionally, although proposition 3.1 settles the accuracy question when data are bandlimited, it remains to be studied how faithful the curvelet transform can be in the presence of *aliased data*. Aliasing occurs when, for example, a function with a discontinuity is discretized by pointwise evaluation. In image processing this typically happens in the neighborhood of an edge. Yet not all hope is lost, because of geometric regularity along the edge. A complete theory of approximation for curvelets (or wavelets for that matter) needs to solve this sampling issue.

Finally, the architecture of our wave atom solver can probably be improved in a variety of ways. For example, predicting the values of the large curvelet/wave atom matrix elements in some way involving geometrical optics is a natural idea. We have in mind a *parametrix* construction, as in [74], coupled with the Phase-Flow Method for solving the Hamiltonian ODE system.

5.3 Thinking Outside the Grid

Trying to fit the "correct" applied harmonic analysis tool to a numerical analysis problem is sometimes a discouraging experience – wavelet enthusiasts should be warned – but it has the merit of offering its own *intellectual challenges*. We believe a research project is all the more interesting if it ends up somewhere else than it was intended to. In our case:

- The Phase-Flow Method, developed by Lexing Ying and Emmanuel Candès in [90], is a very clean answer to the problem of computing the position of shifted diagonals in the curvelet matrix of wave propagators in optimal complexity. Without the underlying motivation their project may never have seen the light of day.

- Developing a fast curvelet transform was an imperative prerequisite at the time we started doing numerical experiments on wave equations. The code has since then been made available (http://www.curvelet.org) and is now used in a variety of unexpected contexts, mostly for inverse problems in seismic imaging involving denoising and compression of 'curvelet-looking' bandlimited wavefronts. For more info, see for example Felix Herrmann's webpage http://slim.eos.ubc.ca.

- Creating a wave atom transform as an alternative to curvelets had been our next milestone. It turns out that their construction raises some fundamental questions in wavelet theory and filterbank architecture, if for example one wishes to make wave atoms compactly supported in space.

- Our efforts to input some ideas from high-dimensional numerical analysis into the wave atom solver, and the resulting interesting mathematics, is another testimony to the challenges raised by the implementation of a sparsity-only method.

As we speak, information theory is being redefined by unorthodox ideas as part of the quest for ideal data representation [23]. It is our hope that questioning conventional wisdom could shape new research directions in scientific computing as well.

Appendix A

Additional Proofs for Chapter 2

A.1 Additional Proofs for Section 2.2

Proof of proposition 2.1. These four properties were already formulated in [74], although with a slightly weaker definition of pseudo-distance. Properties 1 and 2 are not proved in that reference, and property 3 is not extensively documented. We give the justification for these three results for completeness.

Claim (1). We are to show that $d(\mu, \mu') \asymp d(\mu', \mu)$. With $e_\mu = \xi_\mu/|\xi_\mu|$, this is

$$|\langle e_\mu, \Delta x\rangle| + |\Delta x|^2 + |\Delta\theta|^2 \asymp |\langle e_{\mu'}, \Delta x\rangle| + |\Delta x|^2 + |\Delta\theta|^2.$$

It is sufficient to notice that

$$|\langle e_\mu, \Delta x\rangle| + |\Delta x|^2 + |\Delta\theta|^2 \asymp |\langle e_\mu, \Delta x\rangle| + |\langle e_{\mu'}, \Delta x\rangle| + |\Delta x|^2 + |\Delta\theta|^2.$$

In order to justify the nontrivial inequality, use the law of cosines illustrated in Figure A.1:

$$\begin{aligned}
|\langle e_\mu, \Delta x\rangle|^2 + |\langle e_{\mu'}, \Delta x\rangle|^2 &= \sin^2|\Delta\theta|\,(d_\mu^2 + d_{\mu'}^2) \\
&= \sin^2|\Delta\theta|\,|\Delta x|^2 \pm 2|\langle e_\mu, \Delta x\rangle|\,|\langle e_{\mu'}, \Delta x\rangle|\cos|\Delta\theta| \\
&\le \sin^2|\Delta\theta|\,|\Delta x|^2 + 2|\langle e_\mu, \Delta x\rangle|\,|\langle e_{\mu'}, \Delta x\rangle|.
\end{aligned}$$

It follows that $||\langle e_\mu, \Delta x\rangle| - |\langle e_{\mu'}, \Delta x\rangle|| \le C\cdot|\Delta\theta||\Delta x| \le C\cdot(|\Delta\theta|^2 + |\Delta x|^2)$ and, therefore,

$$|\langle e_\mu, \Delta x\rangle| + |\langle e_{\mu'}, \Delta x\rangle| \le C\cdot(2|\langle e_\mu, \Delta x\rangle| + |\Delta\theta|^2 + |\Delta x|^2).$$

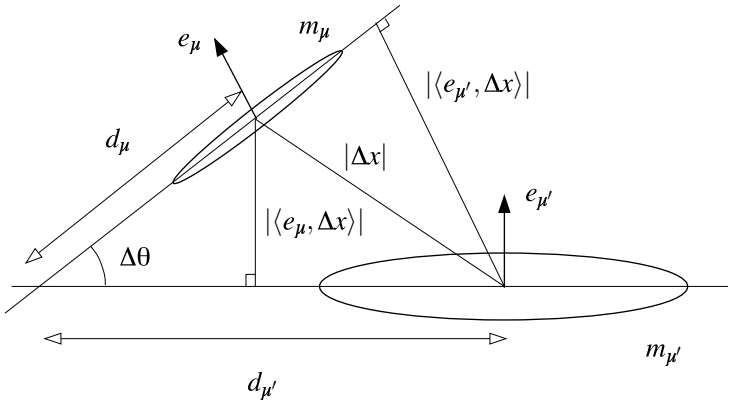

Figure A.1: Relative position and orientation of two curvelet molecules in x-space. The ellipses indicate their essential support.

Claim (2). Recall that $\omega(\mu,\mu') = 2^{|j-j'|}(1+2^{\min(j,j')}d(\mu,\mu'))$. Let us show that $d(\mu,\mu') \le C\cdot(d(\mu,\mu'')+d(\mu'',\mu'))$. To simplify notations, set in the coordinates defined by $\{e_\mu, e_\mu^\perp\}$,

$$\begin{array}{lll} x_\mu = (0,0) & x_{\mu'} = (x_1,x_2) & x_{\mu''} = (y_1,y_2) \\ e_\mu = (1,0) & e_{\mu'} = (\cos\alpha,\sin\alpha) & e_{\mu''} = (\cos\beta,\sin\beta) \\ |\theta_l - \theta_{l''}| = |\beta| & |\theta_{l'} - \theta_{l''}| = |\alpha-\beta| & \end{array}$$

It is enough to show that there exists $\epsilon > 0$ such that

$$\begin{aligned} \epsilon|x_1| \le |y_1| + |\cos\alpha(x_1-y_1) + \sin\beta(x_2-y_2)| \\ + (|\beta|+|\alpha-\beta|)(|y_1|+|x_1-y_1|+|y_2|+|x_2-y_2|), \end{aligned}$$

because then $(|\beta|+|\alpha-\beta|)(|y_1|+|x_1-y_1|+|y_2|+|x_2-y_2|) \le C\cdot(|\beta|^2+|\alpha-\beta|^2+|y_1|^2+|x_1-y_1|^2+|y_2|^2+|x_2-y_2|^2)$. By contradiction let us assume that the inequality fails. Then we must have $|y_1| < \epsilon|x_1|$. It is always true that $|x_1-y_1|+|y_1| \ge |x_1|$ so it is necessary that $|\beta|+|\alpha-\beta| < \epsilon$. But then $|\alpha| < 2\epsilon$ thus $\cos\alpha > 1-4\epsilon^2$ and $|\sin\alpha| < 2\epsilon$. The term $|\cos\alpha(x_1-y_1)+\sin\beta(x_2-y_2)|$ is therefore always greater than $(1-4\epsilon^2)|x_1-y_1| - \epsilon|x_2-y_2|$. But this quantity must also be less than $\epsilon|x_1-y_1|$, otherwise its sum with $|y_1|$ would exceed $\epsilon|x_1|$. So we must have $|x_2-y_2| > \frac{1-\epsilon-4\epsilon^2}{\epsilon}|x_1-y_1|$. But then

the sum $|y_1| + |x_1 - y_1| + |x_2 - y_2|$ must dominate $\frac{|x_1|}{2\epsilon}$, which implies $|\beta| + |\alpha - \beta| \leq 2\epsilon^2$. By induction, $\alpha = \beta = 0$ and $|y_1| + |x_1 - y_1| \geq |x_1|$ yields a contradiction.

Claim (3). We need to establish that $\sum_{\mu_1} \omega(\mu_0, \mu_1)^{-N} \cdot \omega(\mu_1, \mu_2)^{-N} \leq C_N \cdot \omega(\mu_0, \mu_2)^{-(N-1)}$. We closely follow and expand the argument in [74]. We will need to use $d(\mu_0, \mu_1) \asymp d(\mu_1, \mu_0)$, as we have just showed. Define I_{μ_1} by

$$\begin{aligned} I_{\mu_1} &:= \omega(\mu_2, \mu_1)^{-N} \cdot \omega(\mu_1, \mu_0)^{-N} \\ &= \Big(2^{|j_2 - j_1| + |j_1 - j_0|}(1 + 2^{\min(j_2, j_1)} d(\mu_2, \mu_1))(1 + 2^{\min(j_0, j_1)} d(\mu_0, \mu_1))\Big)^{-N}. \end{aligned}$$

To ease notations, put temporarily $a_0 = 2^{\min(j_0, j_1)}$, $a_2 = 2^{\min(j_2, j_1)}$, $d_{01} = d(\mu_0, \mu_1)$, and $d_{12} = d(\mu_2, \mu_1)$. We develop a lower bound on $(1 + a_2 d_{12})(1 + a_0 d_{01}) = 1 + a_2 d_{12} + a_0 d_{01} + a_2 d_{12} a_0 d_{01}$. We make three simple observations: first,

$$a_2 d_{12} + a_0 d_{01} \geq \min(a_2, a_0)(d_{12} + d_{01}) = A_0, \quad \text{and} \quad d_{12} + d_{01} \geq C \cdot d(\mu_0, \mu_2);$$

second,

$$a_2 d_{12} + a_0 d_{01} \geq \max(a_2 d_{12}, a_0 d_{01}) \geq \max(a_2, a_0) \min(d_{12}, d_{01}) = B_0;$$

and third

$$\begin{aligned} a_2 d_{12} a_0 d_{01} &= \max(a_2, a_0) \min(a_2, a_0) \max(d_{12}, d_{01}) \min(d_{12}, d_{01}) \\ &\geq \max(a_2, a_0) \min(a_2, a_0) \min(d_{12}, d_{01}) \frac{d_{12} + d_{01}}{2} = A_0 B_0 / 2. \end{aligned}$$

This gives

$$1 + a_2 d_{12} + a_0 d_{01} + a_2 d_{12} a_0 d_{01} \geq \frac{1}{2}(1 + A_0 + B_0 + A_0 B_0) \geq \frac{1}{2}(1 + A_0)(1 + B_0).$$

We replace the values of A_0, B_0 by their expression, use the relation $A_0 \geq d(\mu_0, \mu_2)$ and obtain

$$I_{\mu_1} \leq C \cdot 2^{-(|j_2 - j_1| + |j_0 - j_1|)N} \cdot \Big(1 + 2^{\min(j_2, j_0, j_1)} d(\mu_2, \mu_0)\Big)^{-N} \cdot (L_1)^{-N} \tag{A.1}$$

with

$$L_1 = 1 + \max(2^{\min(j_2, j_1)}, 2^{\min(j_0, j_1)}) \min(d_{01}, d_{12}).$$

Note that

$$\begin{aligned} L_1 &= \min\Big(1+\max(2^{\min(j_2,j_1)},2^{\min(j_0,j_1)})d_{12}, 1+\max(2^{\min(j_2,j_1)},2^{\min(j_0,j_1)})d_{01}\Big) \\ &\geq \min\Big(1+2^{\min(j_2,j_1)}d_{12}, 1+2^{\min(j_0,j_1)}d_{01}\Big) \end{aligned}$$

and, therefore,

$$\begin{aligned} (L_1)^{-N} &\leq \max\Big((1+2^{\min(j_2,j_1)}d_{12})^{-N}, (1+2^{\min(j_0,j_1)}d_{01})^{-N}\Big) \\ &\leq (1+2^{\min(j_2,j_1)}d_{12})^{-N} + (1+2^{\min(j_0,j_1)}d_{01})^{-N}. \end{aligned}$$

In the sequel we will repeatedly make use of the bound

$$\sum_{k,\ell}(1+2^q d(\mu,\mu'))^{-N} \leq C\cdot 2^{2(j-q)_+}, \tag{A.2}$$

valid for $N \geq 2$, any real q and where the subscript $+$ denotes the positive part. This is justified as follows. Without loss of generality, assume that $\mu' = (j',0,0)$ so that the curvelet $\gamma_{\mu'}$ is nearly vertical and centered near the origin. We recall that $\Delta\theta = \pi\cdot\ell\cdot 2^{-\lfloor j/2\rfloor}$, $\ell = 0,1,\ldots,2^{\lfloor j/2\rfloor}-1$, and $x_\mu = R_{\theta_\mu}D_j^{-1}k$, say. Then the left-hand side is

$$\sum_{\ell=0}^{2^{\lfloor j/2\rfloor}-1}\sum_{k\in\mathbb{Z}^2}\Big(1+2^q(|2^{-j/2}\ell|^2+|2^{-j/2}k_2|^2+|2^{-j}k_1|)\Big)^{-N}. \tag{A.3}$$

For $j \geq q$ this can be seen as a Riemann sum and bounded—up to a numerical multiplicative constant—by the corresponding integral

$$\int_{\mathbb{R}^2}\frac{dx}{2^{-3j/2}}\int_{\mathbb{R}}\frac{dy}{2^{-j/2}}[1+2^q(y^2+x_2^2+|x_1|)]^{-N}$$

which in turn is less than $C\cdot 2^{2(j-q)}$ provided $N \geq 2$. For $j \leq q$, the sum (A.3) essentially consists of a few terms, giving a $O(1)$ contribution. This gives the bound $C\cdot 2^{2(j-q)_+}$.

By symmetry, we can now assume $j_0 \leq j_2$. Let us consider three cases.

- $\mathbf{0 \leq j_2 \leq j_1}$. In that case we have the bound

$$(L_1)^{-N} \leq C\cdot[(1+2^{j_2}d_{01})^{-N} + (1+2^{j_2}d_{12})^{-N}].$$

Summing this quantity over k_1 and ℓ_1 i.e., over all μ_1 that correspond to a given j_1, and using (A.2), we obtain for $j_1 \geq j_2$

$$\begin{aligned}\sum_{\mu_1} I_{\mu_1} &\leq C \cdot (1 + 2^{j_0} d_{02})^{-N} \sum_{j_1 \geq j_2} 2^{-(2j_1 - j_0 - j_2)N} \cdot 2^{2(j_1 - j_2)} \\ &\leq C \cdot 2^{-(j_2 - j_0)N} (1 + 2^{j_0} d_{02})^{-N} = C \cdot \omega(\mu_0, \mu_2)^{-N}.\end{aligned}$$

- $\mathbf{0 \leq j_1 \leq j_0}$. We now have

$$(L_1)^{-N} \leq C \cdot [(1 + 2^{j_1} d_{01})^{-N} + (1 + 2^{j_1} d_{12})^{-N}].$$

According to (A.2), the sum over k_1 and ℓ_1 of $(L_1)^{-N}$ is bounded by a constant independent of j_1. The remaining sum is

$$\sum_{\mu_1} I_{\mu_1} \leq C \cdot 2^{-(j_0 + j_2)N} \sum_{j_1 \leq j_0} 2^{2j_1 N} \cdot (1 + 2^{j_1} d_{02})^{-N}.$$

Observe that $2^{j_1 N}(1 + 2^{j_1} d_{02})^{-N} \leq 2^{j_0 N}(1 + 2^{j_0} d_{02})^{-N}$, therefore

$$\sum_{\mu_1} I_{\mu_1} \leq C \cdot 2^{-(j_2 - j_0)N} (1 + 2^{j_0} d_{02})^{-N} = C \cdot \omega(\mu_0, \mu_2)^{-N}.$$

- $\mathbf{j_0 \leq j_1 \leq j_2}$. In that case we still have

$$(L_1)^{-N} \leq C \cdot [(1 + 2^{j_1} d_{01})^{-N} + (1 + 2^{j_1} d_{12})^{-N}].$$

summed over k_1 and ℓ_1 into a $O(1)$ contribution. What remains is

$$\begin{aligned}\sum_{\mu_1} I_{\mu_1} &\leq C \cdot 2^{-(j_2 - j_0)N} (1 + 2^{j_0} d_{02})^{-N} \sum_{j_0 \leq j_1 \leq j_2} 1 \\ &\leq C \cdot \omega(\mu_0, \mu_2)^{-(N-1)}.\end{aligned}$$

We conclude by collecting the estimates corresponding to the three different cases. Remark that the loss of one (fractional) power of ω in the third case is unavoidable unless one modifies its definition in the spirit of [74]. This would however make notations unnecessarily heavy.

Claim (4). See [74] p. 804. □

Proof of the inequality (2.21). Assume without loss of generality that $\mu = \mu_0$. We may express $S_{\mu'}(\xi)$ as $S_{\mu'_0}(R_{\Delta\theta}\xi)$, with $\Delta\theta = \theta_\mu - \theta_{\mu'}$. We begin by expressing the integral in polar coordinates,

$$\begin{aligned} \xi_1 &= r\cos\theta & (R_{\Delta\theta}\xi)_1 &= r\cos(\theta+\Delta\theta), \\ \xi_2 &= r\sin\theta & (R_{\Delta\theta}\xi)_2 &= r\sin(\theta+\Delta\theta). \end{aligned}$$

As we can see, the cosine factor is not crucial and we may just as well drop it. Consequently,

$$\begin{aligned} \int |S_\mu(\xi)\, S_{\mu'}(\xi)|^n \, d\xi \le C\cdot \int_0^\infty r dr \, \frac{1}{[1+2^{-j}r]^N} \frac{1}{[1+2^{-j'}r]^N} \\ \times \int_0^{2\pi} d\theta [1+a|\sin\theta|]^{-N} [1+a'|\sin(\theta+\Delta\theta)|]^{-N}, \end{aligned}$$

where $a = \frac{2^{-j/2}r}{1+2^{-j}r}$ and $a' = \frac{2^{-j'/2}r}{1+2^{-j'}r}$. This decoupling makes the problem of bounding the inner integral on the variable θ tractable. For example when $a > a' > 1$, following [65] p.56,

$$\int_{-\infty}^{\infty} d\theta [1+a|\theta|]^{-N}[1+a'|\theta+\Delta\theta|]^{-N} \le C\cdot \frac{1}{a}\frac{1}{[1+a'|\Delta\theta|]^N}.$$

We get other estimates for other values and orderings of a and a'. The integral on r is then broken up into several pieces according to the values of a, a', j and j'. It is straightforward to show that each of these contributions satisfies the inequality (2.21). □

A.2 Additional Proofs for Section 2.5

Proof of lemma 2.7. By definition $a^{(\mu)}(x) = 2^{-3j/4} m_\mu \left(D_{2^{-j}} R_{\theta_\mu} x - k\right)$ and, therefore,

$$\begin{aligned} a^{(\mu)}(x) &= \frac{1}{\alpha_\mu}\int (Rf)(a,\theta,b) a^{3/4} 2^{-3j/4} \psi(D_a R_\theta (R_{\theta_\mu}^{-1} D_{2^j}(x+k) - b))\, d\mu \\ &= \frac{1}{\alpha_\mu}\int (Rf)(a,\theta,b)|A|^{1/2} \psi\left(A(x-(\beta-k))\right)\, d\mu, \end{aligned} \tag{A.4}$$

where $A = D_a R_\delta D_{2^j}$ with $\delta = \theta - \theta_\mu$ and $\beta = D_{2^{-j}} R_{\theta_\mu} b$.

Let us first verify the assertion about the support of $a^{(\mu)}$. Recall that over a cell Q_μ, $\beta \in [k_1, k_1+1)\times[k_2,k_2+1)$, and hence for all $b \in Q_\mu$, we have

$$\operatorname{Supp}\psi\left(A(x-(\beta-k))\right) \subset \operatorname{Supp}\psi(Ax) + [0,1]^2.$$

Next Supp $\psi(Ax) \subset A^{-1}[0,1]^2$ with $A^{-1} = D_{2^{-j}} R_{-\delta} D_a^{-1}$. It is not difficult to check that $A^{-1}[0,1]^2 \subset [c_1, c_2) \times [d_1, d_2)$ which then gives (2.65).

There are several ways to prove the property about nearly vanishing moments. A possibility is to show that the Fourier transform of $a^{(\mu)}$ is appropriately small in a neighborhood of the axis $\xi_1 = 0$. We choose a more direct strategy and show that

$$\left| \int \psi(A(x-\beta)) x_1^k \, dx_1 \right| \leq C_m \cdot 2^{-j(m+1)}. \tag{A.5}$$

uniformly over the $(a, \theta, b) \in Q_\mu$. The property (2.66) follows from this fact. Indeed,

$$\int a^{(\mu)}(x_1, x_2) \, x_1^k \, dx_1 = \frac{1}{\alpha_\mu} \int_{Q_\mu} Rf(a, \theta, b) d\mu \int |A|^{1/2} \psi(A(x-\beta)) x_1^k \, dx_1,$$

and the Cauchy-Schwarz inequality gives

$$\begin{aligned} \left| \int a^{(\mu)}(x_1, x_2) \, x_1^k \, dx_1 \right| &\leq \frac{1}{\alpha_\mu} \|Rf\|_{L_2(Q_\mu)} \left(\int_{Q_\mu} \left| \int |A|^{1/2} \psi(A(x-\beta)) x_1^k \, dx_1 \right|^2 d\mu \right)^{1/2} \\ &= \left(\int_{Q_\mu} \left| \int |A|^{1/2} \psi(A(x-\beta)) x_1^k \, dx_1 \right|^2 d\mu \right)^{1/2}. \end{aligned}$$

The uniform bound (A.5) together with the fact that $\int_{Q_\mu} d\mu$ is either 3π or $3\pi/2$ gives (2.66).

We then need to establish (A.5). Let ∂_2 be $\partial/\partial x_2$, recall that by assumptions (2.58) and (2.59), we have that for all $x_2 \in \mathbb{R}$,

$$\int \partial_2^n \psi(x_1, x_2) x_1^k \, dx_1 = 0, \quad k = 0, 1, \ldots, R,$$

and more generally, for each $\alpha \neq 0$ and β

$$\int \partial_2^n \psi(\alpha x_1 + \beta, x_2) x_1^k \, dx_1 = 0, \quad k = 0, 1, \ldots, R. \tag{A.6}$$

We shall use (A.6) to prove (A.5). Letting

$$A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}$$

and with the same notations as before, a simple calculation shows that $a_{21} = -\frac{2^{-j}\sin\delta}{\sqrt{a}}$. As $a \geq 2^{-(j+1)}$ and $|\delta| \leq \pi/2 \cdot 2^{-\lfloor j/2 \rfloor}$, we have

$$|a_{21}| \leq c \cdot 2^{-j}. \tag{A.7}$$

We then write

$$\begin{aligned} \psi(Ax) &= \psi(a_{11}x_1 + a_{12}x_2, a_{21}x_1 + a_{22}x_2) \\ &= \sum_{n=0}^{N-1} D^n\psi(a_{11}x_1 + a_{12}x_2, a_{22}x_2)\frac{(a_{21}x_1)^n}{n!} + O((a_{21}x_1)^N) \end{aligned}$$

and, therefore,

$$\int \psi(Ax)x_1^k\,dx_1 = \sum_{n=0}^{N-1} \frac{a_{21}^n}{n!}\int D^n\psi(a_{11}x_1 + a_{12}x_2, a_{22}x_2)x_1^{n+k}\,dx_1 + O(a_{21}^N)$$

Fix $k \leq D$ and pick $N = D - k + 1$ so that for $n = 0, 1, \ldots, N-1$, $n + k \leq D$. By virtue of (A.6) all the integrals in the sum vanish and the only remaining term is $O(a_{21}^N)$ which because of (A.7) is $O(2^{-jN})$. As a consequence, setting $m = D/2$, we conclude that

$$\left|\int \psi(Ax)x_1^k\,dx_1\right| \leq C_m \cdot 2^{-j(m+1)}, \quad k = 0, 1, \ldots, m;$$

this is the content of (A.5).

The careful reader will notice that inequality (A.5) or equivalently (2.66) is a weaker statement than inequality (2.13) for the definition of nearly vanishing moments. There is no doubt that the stronger estimate (2.13) also holds for curvelet atoms. The proof of this fact uses standard arguments and we choose not to reproduce it here.

Last, the regularity property is a simple consequence of the Cauchy Schwarz inequality;

$$\begin{aligned} \left|a^{(\mu)}(x_1, x_2)\right| &\leq \frac{1}{\alpha_\mu}\int |Rf(a,\theta,b)||A|^{1/2}\|\psi\|_{L_\infty}\,d\mu \\ &\leq \|\psi\|_{L_\infty} \cdot \frac{1}{\alpha_\mu}\|Rf\|_{L_2(Q_\mu)} \cdot \left(\int_{Q_\mu} |A|d\mu\right)^{1/2} \\ &= 2\sqrt{3\pi} \cdot \|\psi\|_{L_\infty}. \end{aligned}$$

These last inequalities used the facts that $|A| \leq 4$ for $(a, \theta, b) \in Q_\mu$ and $\int_{Q_\mu} d\mu \leq 3\pi$. Esti-

mates for higher derivatives are obtained in exactly the same fashion–after differentiation of the integrand. This finishes the proof of the lemma. □

Proof of (2.68). Recall that

$$\alpha_{\mu\mu'} = \left(\int_{Q_{\mu'}} |\mathcal{R}(q(D)\gamma_\mu)(a,b,\theta)|^2 \, d\mu \right)^{1/2}.$$

The first thing to notice is that $q(D)\gamma_\mu$ is still a family of curvelet molecules, because $q(\xi)$ is a multiplier of order zero. Since $\psi_{a,\theta,b}$ also obeys the molecule properties, lemma 2.1 implies the corresponding almost-orthogonality condition. Integrating over $Q_{\mu'}$ does not compromise this estimate, as can be seen by applying the Cauchy-Schwarz inequality. □

Proof of inequality (2.52). Derivatives of $\hat{\gamma}_\mu$ and σ are treated using the following estimates.

$$\begin{aligned} &|\partial_\xi^\alpha \hat{\gamma}_\mu(\xi)| \le C_\alpha \cdot 2^{-3j/4} 2^{-\alpha_1 j} 2^{-\alpha_2 j/2} \\ &|\partial_\xi^\alpha \sigma(\phi^{-1}(x),\xi)| \le C_\alpha \cdot 2^{-|\alpha| j} \qquad \text{on } W_\mu = \mathrm{supp}(\hat{\gamma}_\mu). \end{aligned}$$

We now develop size estimates for the phase perturbation δ. Following closely the discussion in [79], p.407, we claim that on W_μ,

$$|\partial_\xi^\alpha \partial_x^\beta \delta(x,\xi)| \le C_{\alpha\beta} \cdot 2^{-\alpha_1 j} 2^{-\alpha_2 j/2}. \tag{A.8}$$

The derivations in x add no complications. Hence, assume that $\beta = 0$. As the above result (A.8) relies upon the homogeneity of the phase with respect to ξ, we recall a few useful facts about homogeneous functions of degree one:

$\Phi = \Phi_\xi \cdot \xi$ (Euler's theorem),

$\Phi_{\xi\xi} \cdot \xi = 0$ (differentiate the above relation),

$\partial_\xi^\alpha \Phi = O(|\xi|^{1-|\alpha|})$.

It follows from the definition that $\delta(x,\xi_1,0) = 0$ and likewise $\frac{\partial \delta}{\partial \xi_2}(x,\xi_1,0) = 0$. Thus for every n, $\frac{\partial^n \delta}{\partial \xi_1^n}(x,\xi_1,0) = 0$ and $\frac{\partial}{\partial \xi_2}\frac{\partial^n \delta}{\partial \xi_1^n}(x,\xi_1,0) = 0$. Recall that the support conditions are

$|\xi_1| \le C \cdot 2^j$ and $|\xi_2| \le C \cdot 2^{j/2}$. Taylor series expansions about $\xi_2 = 0$ together with homogeneity assumptions give

$$\begin{aligned}
\frac{\partial^{\alpha_1}}{\partial \xi_1^{\alpha_1}} \delta(x,\xi) &= O(|\xi_2|^2 |\xi|^{-1-\alpha_1}) = O(2^{-\alpha_1 j}), \\
\frac{\partial}{\partial \xi_2} \frac{\partial^{\alpha_1}}{\partial \xi_1^{\alpha_1}} \delta(x,\xi) &= O(|\xi_2||\xi|^{-1-\alpha_1}) = O(2^{-j/2} 2^{-\alpha_1 j}), \\
\frac{\partial^{\alpha_2}}{\partial \xi_2^{\alpha_2}} \frac{\partial^{\alpha_1}}{\partial \xi_1^{\alpha_1}} \delta(x,\xi) &= O(|\xi|^{1-\alpha_1-\alpha_2}) = O(2^{-\alpha_1 j} 2^{-\alpha_2 j/2}) \qquad \text{when } \alpha_2 \ge 2,
\end{aligned}$$

as claimed. The point about these estimates is that they exhibit exactly the parabolic scaling of curvelets. We conclude

$$|\partial_\xi^\alpha e^{i\delta(\phi^{-1}(x),\xi)}| \le C_\alpha \cdot 2^{-\alpha_1 j} 2^{-\alpha_2 j/2} \qquad \text{on } W_\mu$$

and therefore (2.52). □

Appendix B

Additional Proofs for Chapter 4

Proof of inequality (4.39). In what follows the notation *sup* refers to the supremum taken over all $x \in B_{x_k}(r_k)$, and over all components of vector or matrix arguments. Put $r_k = r(x_k, j)$. We need to show that

$$\sup |r(x,j) - r_k| \leq \sup \left| \frac{\delta r}{\delta |\nabla g|} \right| \sup |\nabla g(x) - \nabla g(x_k)| \leq C \cdot r_k.$$

On the one hand,

$$\frac{\delta r}{\delta |\nabla g|} = -\frac{1}{2} \frac{2^{2j}}{(2^j + 2^{2j}|\nabla g(x)|)^{3/2}} = -\frac{1}{2} 2^{2j} r_k^3. \tag{B.1}$$

On the other hand,

$$|\nabla g(x) - \nabla g(x_k)| \leq r_k \sup |Hg(x)|, \tag{B.2}$$

where $Hg(x)$ is the Hessian of g. In order to estimate $|Hg(x)|$, recall Landau's inequality for the interval $[0,1]$ which reads

$$\|f'\|_\infty \leq \frac{2}{h} \|f\|_\infty + \frac{h}{4} \|f''\|_\infty,$$

for all $0 \leq h \leq 1$ (see for example [3]). This inequality needs to be extended to two dimensions and applied twice with $\partial^\alpha g$ in place of f, $\alpha = (1,0)$ and $(0,1)$ respectively (where g is understood to be adequately extended to zero outside of $B_{x_k}(r_k)$). Upon choosing $h = C \cdot \|\partial^\alpha g\|_\infty^{1/2} \leq C \cdot 2^{-j} r_k^{-1} \leq 1$, with the constant C determined by the condition $h \leq 1$, it follows that (the sup is still over $x \in B_{x_k}(r_k)$),

$$\sup |Hg(x)| \leq C \cdot \sup |\nabla g(x)|^{1/2} \leq C \cdot 2^{-j} r_k^{-1}. \tag{B.3}$$

As always the constant C changes from line to line. From equations (B.1), (B.2) and (B.3), we check that

$$\sup |r(x,j) - r_k| \leq C \cdot 2^j r_k^3$$

This is dominated by $C \cdot r_k$ because $r_k \leq 2^{-j/2}$, and we are done.

Proof of lemma 4.6. The coarea for BV functions in the unit square $\Omega \subset \mathbb{R}^2$ is

$$\int |\nabla g(x)|\, dx = \int_{-\infty}^{\infty} H^1(g^{-1}(t) \cap \Omega)\, dt. \tag{B.4}$$

Written as above, the formula is valid for Lipschitz functions, the quantity $|\nabla g(x)|$ must be interpreted in a suitable measure-theoretic sense and the proof is rather technical. For C^2 functions, the proof is more accessible and can be found in [92], pp. 76 and following.

In our case $g \in C^2([0,1]^2)$ so the level sets $g^{-1}(t) \cap \Omega$ of g have bounded Hausdorff-H^1 measure for almost every t, and

$$g^{-1}(t) \cap \Omega \equiv \partial N_t, \qquad \text{a.e. } t,$$

where $N_t = \{x \in \Omega, g(x) \leq t\}$. We can let $X_t = N_t \backslash N_{-t}$ as in the wording, and apply the coarea formula to the function defined as

$$\tilde{g}(x) = \begin{cases} g(x) \text{ if } |g(x)| \leq t, \\ -t \text{ if } g(x) < -t, \\ t \text{ if } g(x) > t. \end{cases}$$

Since sets of zero measure do not contribute in the integral in t, we obtain

$$\int_{X_t} |\nabla g(x)|\, dx = \int_{-t}^{t} H^1(\partial N_t)\, dt \leq 2t \sup_u H^1(\partial N_u).$$

We leave it as an exercise to the interested reader to prove that there is another, perhaps more visual way to derive the above formula from the Reynolds transport theorem.

Proof of inequality (4.44).

Let $\tilde{\epsilon} > 0$ and $\tilde{r}_k(x', \xi', t)$ be the separation remainder of $b_k(x', \xi', t)$ for that $\tilde{\epsilon}$ in L^2. We

invoke the strong version of Lemma 4.2 to obtain control on $\tilde{r}_k$ in $W^{s,\infty}$,

$$|\partial_{\xi'}^{\alpha}\partial_{x'}^{\beta}\tilde{r}_k(x',\xi',t)| \leq C_{\alpha\beta}\tilde{\epsilon}.$$

In the original variables x and ξ, let $r_k(x,\xi,t) = \tilde{r}_k(x',\xi',t)$ so the condition becomes

$$|\partial_{\xi}^{\alpha}\partial_{x}^{\beta}r_k(x,\xi,t)| \leq C_{\alpha\beta}\tilde{\epsilon}(1+|\xi|)^{-|\alpha|/2+|\beta|/4},$$

i.e., $\frac{r_k}{\tilde{\epsilon}}$ is a symbol of order zero and type $(1/2, 1/4)$. Owing to the decomposition $K_{\text{stat}} = \sum_k K_k$ using indicators $q_k(x)$, the total separation remainder is actually the sum $r = \sum_k r_k$. Although each sum contains $O(2^j)$ terms, by the constant overlap property (4.39) for each given x there is a constant number of terms (independent of j) contributing in $\sum_k K_k$. Likewise, the separated components of $q_k(x)$ are all supported on balls centered at x_k with radius twice the diameter of $\text{supp}(q_k)$, so for each given x there is a constant number of terms contributing in $\sum_k r_k$. Hence the symbol property transfers to r,

$$|\partial_{\xi}^{\alpha}\partial_{x}^{\beta}r(x,\xi,t)| \leq C_{\alpha\beta}\tilde{\epsilon}(1+|\xi|)^{-|\alpha|/2+|\beta|/4}.$$

We conclude by standard pseudo-differential calculus that r is bounded in L^2 with a norm not exceeding $C\tilde{\epsilon}$ for some constant C, which by choosing $\tilde{\epsilon}$ small enough can be made less than $\epsilon/4$ as in equation (4.44). The point of the analysis is that the L^2 bound on r is not only small but independent of j.

Bibliography

[1] J. P. Antoine, R. Murenzi, Two-dimensional directional wavelets and the scale-angle representation. *Sig. Process.* **52** (1996), 259–281.

[2] A. Averbuch, L. Braverman, R. Coifman, M. Israeli, A. Sidi, Efficient computation of oscillatory integrals via adaptive multiscale local Fourier bases, *Appl. Comput. Harm. Anal.* **9-1** (2002) 19–53.

[3] Yu. V. Babenko, Pointwise inequalities of Landau-Kolmogorov-type for functions defined on a finite segment, *Ukr. Math. J.* **53-2** (2001) 270–275.

[4] E. Bacry, S. Mallat and G. Papanicolaou, A Wavelet Based Space-Time Adaptive Numerical Method for Partial Differential Equations, *Math. Model. and Num. Anal.* **26-7** (1992) 793.

[5] H. Bahouri, J. Y. Chemin, Quasilinear Wave Equations and Microlocal Analysis, *Proc. ICM 2002*, vol. III, 141–153.

[6] G. Beylkin, Imaging of discontinuities in the inverse scattering problem by inversion of a causal generalized Radon transform, *J. Math. Phys.* **26** (1985) 99–108.

[7] G. Beylkin, R. R. Coifman, V. Rokhlin, Fast wavelet transforms and numerical algorithms. *Comm. on Pure and Appl. Math.* **44** (1991), 141–183.

[8] G. Beylkin, M. J. Mohlenkamp, Algorithms for numerical analysis in high dimensions, *SIAM J. Sci. Comput.* **26-6** (2005) 2133–2159.

[9] G. Beylkin and K. Sandberg, Wave propagation using bases for bandlimited functions, *Wave Motion*, **41-3** (2005) 263-291

[10] J. M. Bony, Interaction des singularités pour les équations aux dérivées partielles non-linéaires, *Séminaire Goulaouic-Meyer-Schwartz* (1981-1982) exp. no. 2.

[11] A. Brandt, *Multiscale Scientific Computation: Review 2000*, Weizmann Institute of Science, 2000.

[12] J. Brüning and V. W. Guillemin (editors), *Mathematics Past and Present: Fourier Integral Operators*, selected classical articles by J. J. Duistermaat, V. W. Guillemin and L. Hörmander, Springer, 1994.

[13] E. J. Candès, Harmonic analysis of neural networks. *Appl. Comput. Harmon. Anal.* **6** (1999), 197–218.

[14] E. J. Candès, L. Demanet, Curvelets and Fourier integral operators. *C. R. Acad. Sci. Paris, Ser. I* **336** (2003), 395–398.

[15] E. J. Candès and L. Demanet, The curvelet representation of wave propagators is optimally sparse, *Comm. Pure Appl. Math.* **58-11** (2005) 1472–1528.

[16] E. J. Candès, L. Demanet, D. L. Donoho, L. Ying, Fast Discrete Curvelet Transforms, submitted to *SIAM Mult. Model. Sim.* (2005)

[17] E. J. Candès and D. L. Donoho, Ridgelets: the key to higher-dimensional intermittency? *Phil. Trans. R. Soc. Lond. A.* **357** (1999), 2495–2509.

[18] E. J. Candès, D. L. Donoho, Curvelets—a surprisingly effective nonadaptive representation for objects with edges. *Curves and Surfaces*, 105–120. Edited by C. Rabut, A. Cohen, and L. L. Schumaker, Vanderbilt University Press, Nashville, TN, 2000.

[19] E. J. Candès and D. L. Donoho. Recovering edges in ill-posed inverse problems: Optimality of curvelet frames. *Ann. Statist.* **30** (2002), 784 –842.

[20] E. J. Candès, D. L. Donoho, New tight frames of curvelets and optimal representations of objects with piecewise C^2 singularities. *Comm. on Pure and Appl. Math.* **57** (2004), 219–266.

[21] E. J. Candès, D. L. Donoho, Continuous curvelet transform: I. resolution of the wavefront set, *Appl. Comput. Harmon. Anal.* **19-2** (2005), 162–197.

[22] E. J. Candès, F. Guo, New multiscale transforms, minimum total variation synthesis: Applications to edge-preserving image reconstruction. *Signal Processing* **82** (2002), 1519–1543.

[23] E. J. Candès, J. K. Romberg, and T. Tao, Robust uncertainty principles: Exact recovery from highly incomplete Fourier information, *IEEE Trans. Info. Th.* **52-2** (2006), 489–509.

[24] C. Chaux, L. Duval, J. C. Pesquet, 2D Dual-Tree M-band Wavelet Decomposition, *Proc. ICASSP* (2005)

[25] H. Cheng et al., A wideband fast multipole method for the Helmholtz equation in three dimensions, *J. Comput. Phys.*, to appear, 2006.

[26] A. Cohen, *Numerical Analysis of Wavelet Methods*, North-Holland, Elsevier, 2003.

[27] A. Cohen, W. Dahmen, R. DeVore, Adaptive wavelet methods for elliptic operator equations—Convergence rates, *Math. Comp.* **70-233** (2000) 27–75.

[28] A. Córdoba, C. Fefferman, Wave packets and Fourier integral operators. *Comm. PDE* **3(11)** (1978), 979–1005.

[29] G. David and J. L. Journé, A boundedness criterion for generalized Calderó-Zygmund operators, *Ann. Math.* **120** (1984) 371–397

[30] J. M. Delort, *F.B.I. Transformation, second microlocalization and semilinear caustics*, Springer, LNM 1522, 1991.

[31] M. N. Do and M. Vetterli, Contourlets, in *Beyond Wavelets*, ed. G. V. Welland, Academic Press, 2003.

[32] M. N. Do and M. Vetterli, The contourlet transform: An efficient directional multiresolution image representation, *IEEE Trans. Im. Proc.*, to appear, 2005.

[33] D. L. Donoho and M. R. Duncan. Digital Curvelet Transform: Strategy, Implementation, Experiments. Technical Report, Stanford University, 1999.

[34] M. Dorobantu, B. Engquist, Wavelet-based numerical homogenization, *SIAM J. Num. Anal.* **35-2** (1998) 540–559.

[35] H. Douma and M. V. de Hoop. Wave-character preserving prestack map migration using curvelets. Presentation at the *Society of Exploration Geophysicists*, Denver, CO, 2004.

[36] J. Duistermaat, *Fourier integral operators*. Birkhauser, Boston, 1996.

[37] B. Engquist, S. Osher, S. Zhong, Fast wavelet based algorithms for linear evolution equations, *SIAM J. Sci. Comput.* **15-4** (1994) 755–775.

[38] L. C. Evans, *Partial Differential Equations*, Graduate Studies in Mathematics, vol. 19, AMS, 1998.

[39] C. Fefferman, A note on spherical summation multipliers. *Israel J. Math.* **15** (1973), 44–52.

[40] A. G. Flesia, H. Hel-Or, A. Averbuch, E. J. Candès, R. R. Coifman and D. L. Donoho. Digital implementation of ridgelet packets, *Beyond Wavelets*, ed. J. Stoeckler and G. V. Welland, Academic Press, 2003.

[41] G. Folland, *Harmonic analysis in phase space*, Princeton University Press, 1989.

[42] M. Frazier, B. Jawerth, G. Weiss, *Littlewood-Paley theory and the study of function Spaces*. CBMS, 79, American Mathematical Society, Providence, 1991.

[43] L. Garding, *Singularities in linear wave propagation*. Lecture Notes in Mathematics, 1241, Springer, 1987.

[44] K. Guo, D. Labate, W. Lim, G. Weiss, and E. Wilson, Wavelets with Composite Dilations, *Electr. Res. Ann. AMS* **10** (2004), 78–87.

[45] K. Guo, D. Labate, W. Lim, G. Weiss, and E. Wilson, Wavelets with Composite Dilations and their MRA Properties, *Appl. Comput. Harmon. Anal.*, to appear, 2005.

[46] W. Hackbusch, A sparse matrix arithmetic based on H-matrices. Part I: Introduction to H-matrices. *Computing* **62** (1999) 89–108.

[47] G. Hennenfent and F. J. Herrmann. Seismic denoising with unstructured curvelets, *Comput. in Sci. Eng.*, to appear, 2006.

[48] F. J. Herrmann and E. Verschuur. Separation of primaries and multiples by non-linear estimation in the curvelet domain. In *EAGE 66th Conference & Exhibition Proceedings*, 2004.

[49] F. J. Herrmann, P. P. Moghaddam, C. C Stolk, Sparsity- and continuity-promoting seismic image recovery with curvelet frames, submitted, 2006.

[50] M. V. de Hoop, J. H. le Rousseau and R. Wu, Generlization of the phase-screen approximation for the scattering of acoustic waves, *Wave Motion* **31-1** (2000) 43–70.

[51] L. Hörmander, *The Analysis of Linear Partial Differential Operators*, 4 volumes, Springer, 1985.

[52] A. Iserles, S.P. Norsett, S. Olver, Highly oscillatory quadrature: The story so far. Technical Report NA2005-06, University of Cambridge, 2005.

[53] S. Jaffard, Wavelet methods for fast resolution of elliptic problems, *SIAM J. Num. Anal.* **29-4** (1992) 965–986.

[54] I. Kay, J. B. Keller, Asymptotic Evaluation of the Field at a Caustic, *J. Appl. Phys.* **25** (1954) 876.

[55] E. Kolaczyk, *Wavelet Methods for the Inversion of Certain Homogeneous Linear Operators in the Presence of Noisy Data*, Ph.D. thesis, Stanford University, 1994.

[56] D. Labate, W. Lim, G. Kutyniok and G. Weiss, "Sparse Multidimensional Representation using Shearlets, *SPIE Conf. Wavelets XI*, San Diego, CA, 2005.

[57] P. Lax, Asymptotic solutions of oscillatory initial value problems. *Duke Math J.* **24** (1957), 627–646.

[58] G. Lebeau, Propagation des ondes dans les variétés à coins, *Ann. Sci. ENS* **30** (1997) 429–297.

[59] R. J. LeVeque, *Finite Volume Methods for Hyperbolic Problems*, Cambridge University Press, 2002.

[60] J. L. Lions and E. Magenes, *Non-Homogeneous Boundary Value Problems and Applications*, vol. 1, Springer, 1972.

[61] D. Ludwig, Exact and Asymptotic Solutions of the Cauchy Problem, *Comm. Pure Appl. Math.* **13** (1960) 473–508.

[62] S. Mallat, *A Wavelet Tour of Signal Processing*. Second edition. Academic Press, Orlando-San Diego, 1999.

[63] R. Melrose and N. Ritter, Interaction of nonlinear progressive waves, *Annals of Math.*, **121** (1985) 187–213.

[64] Y. Meyer, *Ondelettes et opérateurs*. Hermann, Paris, 1990.

[65] Y. Meyer, R. R. Coifman, *Wavelets, Calderón-Zygmund and Multilinear Operators*. Cambridge University Press, Cambridge, 1997.

[66] C. Moler, C. Van Loan, Nineteen Dubious Ways to Compute the Exponential of a Matrix, Twenty-Five Years Later, *SIAM Review* **45-1** (2003) 3–49.

[67] R. Murenzi, *Ondelettes Multidimensionelles et Applications à l'Analyse d'Images,*, Ph.D. thesis, Université catholique de Louvain, Louvain-la-Neuve, 1990.

[68] A. Seeger, C. Sogge, E. Stein, Regularity properties of Fourier integral operators. *Annals of Math.* **134** (1991), 231–251.

[69] E. Seré, Localisation fréquentielle des paquets d'ondelettes, *Rev. Mat. Iber.*, **11-2** (1995) 334–354.

[70] E. P. Simoncelli and W T Freeman. The Steerable Pyramid: A Flexible Architecture for Multi-Scale Derivative Computation. *IEEE Second Int'l Conf on Image Processing.* Washington DC, October 1995.

[71] E. P. Simoncelli, W. T. Freeman, E. H. Adelson, and D. J. Heeger. Shiftable multi-scale transforms [or what's wrong with orthonormal wavelets]. *IEEE Trans. Information Theory, Special Issue on Wavelets* **38** (1992), 587–607.

[72] J. Sjostrand, Singularités analytiques microlocales, *Astérisque* **95** (1982).

[73] H. Smith, A Hardy space for Fourier integral operators. *J. Geom. Anal.* **8** (1998), 629–653.

[74] H. Smith, A parametrix construction for wave equations with $C^{1,1}$ coefficients. *Ann. Inst. Fourier (Grenoble)* **48**(1998), 797–835.

[75] C. Sogge, *Fourier Integrals in Classical Analysis*. Cambridge University Press, 1993.

[76] J. L. Starck, E. J. Candès, and D. L. Donoho. The curvelet transform for image denoising. *IEEE Trans. Im. Proc.*, **11-6** (2002), 670–684.

[77] J. L. Starck, N. Aghanim and O. Forni, Detecting cosmological non-Gaussian signatures by multi-scale methods. *Astron. and Astroph.* **416** (2004), 9–17.

[78] J. L. Starck, M. Elad, and D. L. Donoho. Redundant multiscale transforms and their application for morphological component analysis. *Adv. Imag. Elec. Phys.* **132** (2004).

[79] E. Stein, *Harmonic Analysis*. Princeton University Press, Princeton, N.J., 1993.

[80] C. Stolk, *On the Modeling and Inversion of Seismic Data*, Ph.D. thesis, Utrecht University, 2000.

[81] C. Stolk, M. de Hoop, Microlocal analysis of seismic inverse scattering in anisotropic elastic media. *Comm. Pure and Appl. Math.* **55** (2002), 261–301.

[82] W. W. Symes, *Mathematics of Reflection Seismology*, Lecture notes, Rice University, 1995.

[83] T. Tao, The weak-type (1, 1) of Fourier integral operators of order -(n-1)/2, *J. Aust. Math. Soc.* **76** (2004) 1–21.

[84] M. Taylor, Reflection of singularities of solutions to systems of differential equations. *Comm. Pure and Appl. Math.* **28** (1975), 457–478.

[85] F. Trèves, *Introduction to pseudo-differential and Fourier integral operators*. Plenum press, 1982, 2 volumes.

[86] P. Vandergheynst and J. F. Gobbers, Directional dyadic wavelet transforms: Design and algorithms. *IEEE Trans. Im. Proc.* **11-4** (2002), 363–372.

[87] A. Vasy, Propagation of singularities for the wave equation on manifolds with corners, *Annals of Math.*, to appear, 2006.

[88] L. Villemoes, Wavelet packets with uniform time-frequency localization, *Comptes-Rendus Mathematique*, **335-10** (2002) 793–796.

[89] G. Whitham, *Linear and Nonlinear Waves*. Wiley Interscience, 1999.

[90] L. Ying, E. J. Candès, The Phase-Flow Method, *J. Comput. Phys.*, to appear, 2006.

[91] L. Ying, L. Demanet, E. J. Candès, 3D Curvelet Transform, *Proc. Conf. Wavelets XI*, San Diego, 2006.

[92] W. Ziemer, *Weakly Differentiable Functions*, Springer, 1989.