

Validation of FFT-Based Algorithms for Large-Scale Modeling of Wave Propagation in Tissue

John C. Mould, Gregory L. Wojcik, Laura M. Carcione, Makoto Tabei*[‡],
T. Douglas Mast[†], Robert C. Waag*[‡],

Weidlinger Associates, 4410 El Camino Real, Suite 110, Los Altos, CA 94022,

*University of Rochester, [†]Department of Electrical & Computer Engineering,

[‡]Departments of Electrical & Computer Engineering and Radiology, Rochester, NY 14642

[†]Applied Research Laboratory, The Pennsylvania State University, University Park, PA 16801

Abstract - We investigate accuracy of existing 2D pseudospectral and k-space formulations for simulating acoustic propagation in tissue or model scattering media. They are intended to provide insight into tissue-ultrasound interaction and a "test bed" for aberration correction schemes in medical imaging. Both methods employ FFT's to evaluate spatial derivatives to high accuracy on coarse grids. The primary difference lies in the approach to time integration. Scattering in large-scale, 2D, inhomogeneous media is included. We compare simulations against analytical solutions to illustrate spatial and temporal discretization required for acceptable solutions.

INTRODUCTION

The medium is represented by a uniform Cartesian grid where pressure/stiffness and velocity/density are unknowns/parameters at discrete points. Spectral operators in space enable accuracy and computational efficiency in very large models. However, inhomogeneities are often represented as piecewise constant from node to node, rather than smooth. The resulting stairstep can produce spurious diffractions at edges/corners, inaccurate reflections and transmissions at interfaces and local Gibbs phenomena, by approximating derivatives at a material discontinuity. Thus, the efficiency permitted by coarse spectral grids is compromised by the need to resolve interface derivatives.

For example, scattering by a soft cylinder requires only two nodes per wavelength inside and outside the cylinder for accurate propagation, but significantly more nodes per wavelength are necessary to reduce interface artifacts. Interface artifacts are quantified for a single interface, 1D

multilayer models, and cylindrical scatterers. Abdominal wall cross sections with coarse and fine-scale inhomogeneities are used to explore fidelity of wave propagation versus nodes per wavelength and tissue characteristic lengths. We show that the existing tools are useable in 2D.

The pseudospectral method is implemented in the *SpectralFlex* code. *Kbench* implements the k-space method.

PSEUDOSPECTRAL AND K-SPACE METHODS

The pseudospectral and k-space methods were formulated to provide efficient high-accuracy solutions to long range wave propagation problems. In fact, they debuted during the same year [1,2]. We briefly describe the two methods as implemented in [3,4], highlighting the major similarities and differences.

Both use FFT's to evaluate spatial derivatives to high accuracy on coarse grids. The primary difference lies in their respective approaches to time integration. Note that coarse spatial grids provide the primary incentive for FFT based (or any high order) method. The computational burden is linear in the number of timesteps per cycle, for both 2D and 3D. Including the timestep, computational burden is proportional to the number of Points Per Wave (PPW)³ in 2D or (PPW)⁴ in 3D.

SpectralFlex adopts a 4th order staggered Adams Bashforth ABS4 time integrator [5]. Among general purpose integrators, this is close to optimal for the current applications - 2-3 digits of accuracy for a wave propagating several hundred wavelengths on the coarsest possible grid. The stability limit for ABS4 in 2D is CFL = 0.3. The CFL number is defined as: $CFL = \Delta t / (\Delta x / c)$, where Δt is the timestep, c is the wavespeed and Δx is the cell size.

Accuracy frequently requires a smaller timestep, say CFL = 0.1. *Kbench* implements a time integrator in k-space based on the exact solution for waves propagating in a homogeneous medium [4]. It outperforms the general purpose ABS4 time integrator for weak scatterers in a homogeneous linear acoustic medium. ABS4 becomes more efficient when the scattering objects have a larger impedance contrast.

LONG RANGE PROPAGATION

To illustrate the advantages of the FFT based approach for long range propagation, we propagate a 2.5 MHz pulse 200 wavelengths through water using both *SpectralFlex* and *PZFlex*, a finite element code that is second order accurate in both space and time. The center frequency is 2.5 MHz, but spectral content is observable up to 5 MHz. Thus, 4 PPW at 2.5 MHz is the minimum sampling capable of resolving the pulse.

Figure 1 compares exact, *PZFlex* and *SpectralFlex* solutions. *SpectralFlex* used 4 PPW at CFL = 0.1, whereas *PZFlex* used 20 PPW at CFL = 0.8. These discretizations in time and space are typical of those that would be used in real problems. The *SpectralFlex* signal looks good and can be made better by reducing the timestep. The *PZFlex* signal is delayed in time and badly dispersed. A much finer grid is required to achieve reasonable accuracy. Note that at CFL = 1., *PZFlex* becomes a characteristic method, and produces exact results, even at 2 PPW. Unfortunately, this only works for 1D linear problems.

Kbench produces exact results for this example because the time integrator is based on the exact solution for a homogeneous medium.

DISCONTINUITIES

Spectral methods compute highly accurate spatial derivatives of smooth fields. Thus, in homogeneous regions, 2 cells per minimum wavelength (ie, highest spatial frequency) suffice. However, at material interfaces both the pressure and velocity fields should exhibit slope discontinuities as given by (1), where n denotes the normal direction and the

$$\frac{\partial p^+}{\partial n} = \frac{\rho^+}{\rho^-} \frac{\partial p^-}{\partial n} \quad (1)$$

superscript defines the + or - side of the interface. The velocity field likewise exhibits slope discontinuities at interfaces.

Spectral methods enforce smoothness, approximating the jumps in normal derivatives with steep gradients over a few cells. This approximation is quite good at 10-20 cells per wavelength, but less accurate at 2 cells per wavelength. For a staggered grid, as in *SpectralFlex*, the material interfaces coincide with velocity nodes, so we average the density at these interface points. On a regular grid, all the nodes lie away from interfaces, so no averaging is necessary, but the accuracy is even worse than for the staggered grid.




Figure 1. Long range pulse propagation through water.

1D versus exact solutions

Table 1 summarizes material properties used for the 1D benchmarks. Figure 2 illustrates the reflection/transmission of a normally incident pulse at a water/fat interface as modeled by *SpectralFlex*. To plotting accuracy, the transmitted signals appear exact (because it has much larger amplitude than the reflected wave). However, the error in the reflected signal is readily apparent at 4 PPW, and barely visible at 6 PPW.

Figure 3 shows results for a water /bone interface. In this case, errors are visible in both the reflected and transmitted signals at 4 PPW. In both codes, the most pathological case is varying density/constant stiffness. Fortunately, most tissues have a higher contrast in stiffness than density [6], so this worst case is seldom encountered. As shown in Fig. 4 (density=1000, 928 kg/m³) errors in the reflected wave are visible even at 12 PPW.

Table 1 – Material Properties for 1-D benchmarks

Material	Density [kg/m ³]	Wavespeed [m/sec]
Water	1000.	1500.
Fat	928.	1427.
Conn	1100.	1537.
Musl	1041.	1571.
Livr	1050.	1577.




Figure 2. Reflected pulse at a water/fat interface.




Figure 3. Reflected pulse at a water/bone interface.

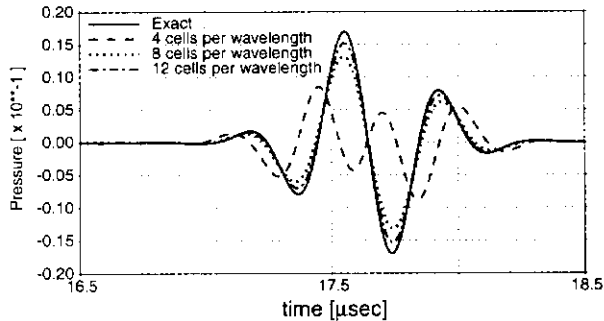


Figure 4. Reflected pulse from worst case interface.

The next benchmark examines propagation through a 1-D approximation of an abdominal cross section. Material parameters are again given in Table 1. Slight errors in the transmitted wave are apparent at 4 PPW (Fig. 5), but not at 8 PPW. Reflected signals (not shown) are similar. Figure 6

illustrates the effect of coarse non-conforming grids. At 4.1 PPW, cell boundaries are misaligned with actual material interfaces by up to ½ cell. This is, of course, the case for any real model with discontinuous material properties. Properties are assigned based on the center of the cell. The errors introduced by this sampling dwarf all others. More will be said about this in a later section.




Figure 5. Pulse transmitted through 1-D abdominal wall model.




Figure 6. Pulse transmitted through 1-D approximation of abdominal wall. Non-conforming grid.

Scattering by cylinders

In addition to the numerical errors at interfaces, approximations are introduced by the stair-step representation of curved surfaces. To quantify these approximations, we consider 3 mm radius fat and bone cylinders immersed in water and insonified by the usual 2.5 MHz pulse. We compute the difference between exact and numerical signals for each timestep at 128 locations at 6 mm radius, and equal spacing in theta. We use the L² norm of this matrix as an error metric. Figure 7a shows the L² error vs PPW for *kbench* and *SpectralFlex* at CFL = 0.2.

The curves are similar, though k_{bench} is slightly more accurate. For the larger contrast bone case in Figure 7b, similar trends are evident, but in this case *SpectralFlex* is more accurate. The error is tending to zero as the PPW increases. The rate of convergence is not quite quadratic. For context, Fig. 12 shows waveforms for L^2 error near 0.01.

Table 2 - Material Properties for Cylinders

Material	Wavespeed [m/sec]	Density [kg/m ³]
Water	1524.	993.
Fat	1478.	950.
Bone	3540.	1990.

Figure 7c illustrates that at low CFL, the error due to time integration tends to zero. For this problem, k_{bench} permits reasonable accuracy at roughly double the *SpectralFlex* timestep. For the bone cylinder, the stability limit of *SpectralFlex* is 0.15 (0.3 in the bone), and k_{bench} can go up to 0.2.

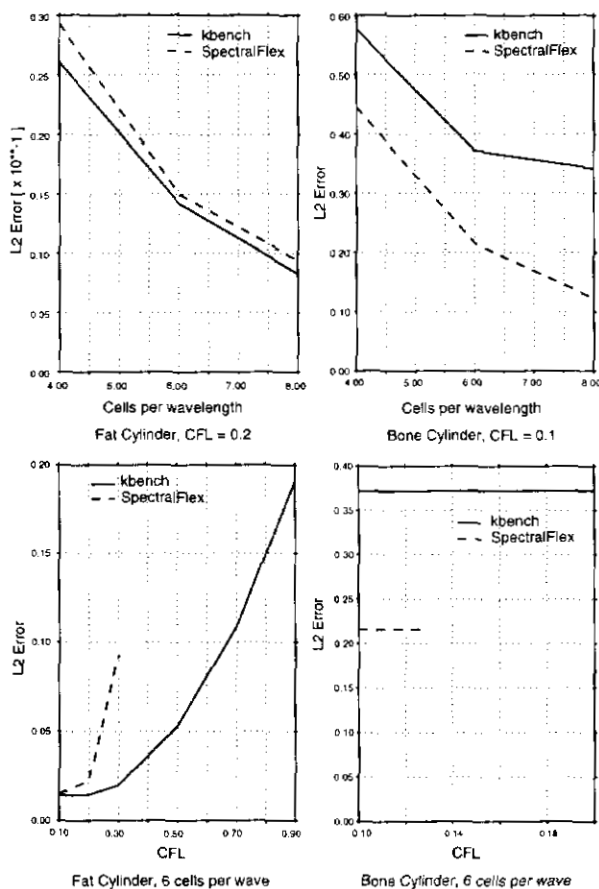


Figure 7. Cylinder benchmarks. Convergence with increasing discretization.




Figure 8. Abdominal wall model.

Tissue examples

Figure 8 shows an abdominal wall cross section [7,8]. This model is insonified by a 4.35 MHz plane wave pulse. Figure 9 displays typical reflected and transmitted signals computed by *SpectralFlex* at 4, 8 and 12 PPW. The grids were defined such that material boundaries always lie in exactly the same place. Again, it is confirmed that even the coarse 4 PPW model produces fairly accurate results.

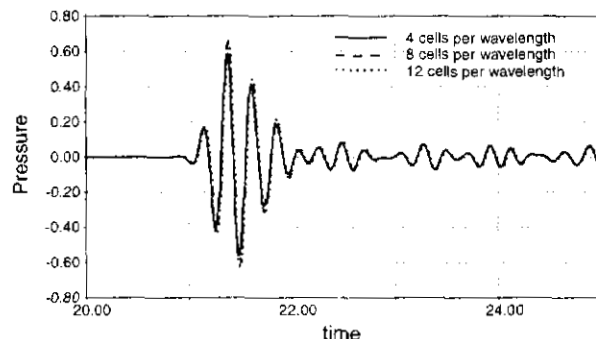


Figure 9. Transmitted pulse from Abdominal wall model.

INTERFACE TREATMENTS

Given that the largest numerical errors in the FFT based methods stem from material interfaces, we look at several interface treatments for reducing those errors.

Jump conditions

One possible method for improving the accuracy at interfaces is to split the solution into smooth and non-smooth parts, and apply the spectral method only to the smooth part. The idea is to introduce

local corrections at material interfaces that enforce the jump conditions exactly. E.g., construct low order polynomials over the cells adjacent to the interface that have zero value and zero slope 1 cell away, and, when added to the continuous part satisfy the jump condition (1) at the interface. Obviously, the correction is not required to be local, but if it covers more than 1 cell, the algorithm will become much more complicated for multiple interfaces. LaVeque [9] discusses such an approach applied to finite difference models.

Figure 10 compares reflected and transmitted signals for coarse models of an interface with and without the jump correction for interface velocity. This example isolates the effects of density changes in that only the density is discontinuous. The bulk modulus is continuous. The correction term improves the computed result, but not to the level of a homogeneous material. A similar correction could be applied to the discontinuity in the velocity gradients. However, it will have a weaker effect on the staggered grid since the leading coefficients are already continuous.

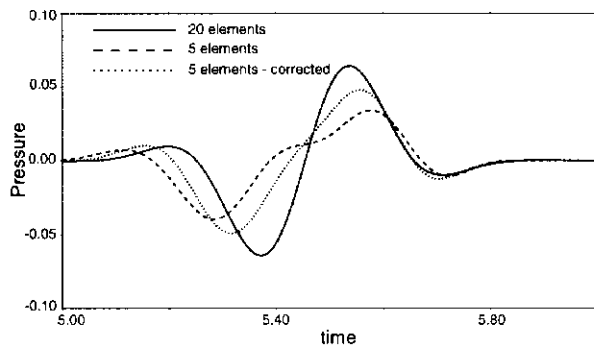



Figure 10. Jump treatment applied to interfaces.

Smoothing (Bandlimitation)

Another approach to improving accuracy at discontinuities is to smooth or bandlimit the model before sampling. This removes unresolvable high spatial frequencies from the model itself. We found that perfect bandlimitation reduced computed signals too much, but a “halfband” filter improves accuracy. The halfband filter is smooth with an amplitude of 0.5 at half the sampling frequency. Figure 11 shows direct and halfband filtered *kbench* models of a 3 mm cylinder using the same number of PPW. The corresponding pressure fields are plotted using a 60 dB bipolar log scale.



a) Unsmoothed b) Halfband Filtered
Figure 11. Direct sampled & bandlimited cylinders. Models and pressure, 60 dB bipolar log scale.

The staircase representation of the cylinder generates diffracted signals at each corner in 11a, but these have disappeared in 11b. Figure 12 shows selected waveforms from the direct and halfband sampled models. The late time diffractions have been removed, and overall L^2 error was reduced from 0.0155 to 0.0105. This exercise demonstrates that smoothing can be useful. However, there are some practical complications. The current procedure computes the smoothed object as the inverse transform of the object’s analytical spectrum multiplied by the filter, and is thus defined only for objects with a known analytical spectrum. The extension to more general models defined on a pixel by pixel level has not yet been demonstrated. Also, continuous variations of material properties produce a large number of distinct materials. In the limit, each cell of the model has different properties. For the purely acoustic case, this presents little difficulty, but when material nonlinearity or viscoacoustic damping is added, the complexity intensifies. E.g., for each wavespeed/damping set, an optimization problem must be solved to compute the appropriate relaxation constants, and these constants must be stored.




Figure 12. Backscattered signals from direct (top) and bandlimited (bottom) models.

Note that this procedure adds information to the model compared to the unsmoothed case. Because smoothing is applied to the analytical cylinder, the continuous variation of material constants provides a richer set of parameters than is available in the unsmoothed representation. As long as the model is known to higher resolution than the grid, information will be added. It is an interesting question whether smoothing would be beneficial on a grid finer than the pixel by pixel model definition. For example, the UOR tissue cross sections [7,8] are the most detailed models we know of. These are represented as piecewise constant with a pixel size 0.085 mm (about 7 PPW for a 2.5 MHz pulse). For a 5 MHz pulse, the coarsest grid would have finer resolution than the model.

Volume averaging of material constants has also been shown effective [10]. This adds additional information compared to the unsmoothed case, and the correction is more local than smoothing. However, the practical difficulties are the same.

As a last resort, increased discretization (brute force) will always converge to an accurate solution. This is a practical solution in 2D, as the above tissue examples indicate.

CONCLUSIONS

Model parameterization is a critical issue and puts all the above results in practical perspective. As shown above, differences in material constants or interface locations cause much larger differences in reflected/transmitted signals than any numerical errors in the FFT based methods. For gaining insight, or as a test-bed for aberration correction schemes, a 4 PPW model is sufficient at frequencies of 2.5 MHz or greater. Fine grids or cell-by-cell representation of material properties are needed only for more accurate rendition of model geometry.

ACKNOWLEDGEMENTS

Sponsored in part (JCM, GLW, LMC) by DARPA and ONR, and monitored by Dr. Wallace A. Smith.

REFERENCES

- [1] N.N. Bojarski, "k-space formulation of the acoustic scattering problem, *J. Acoust. Soc. Am.*, Vol. 53, 374, 1973, presented 84th mtg. ASA, 1972.
- [2] O. Kreiss and J. Olinger, "Comparison of accurate methods for the integration of hyperbolic equations," *Tellus*, vol. 24, 199-215, 1972.
- [3] G. Wojcik, B. Fornberg, R. Waag, L. Carcione, J. Mould, L. Nikodym, T. Driscoll, "Pseudospectral methods for large-scale bioacoustic models," *IEEE Ultrasonics Symp.* 1501-1506, 1997.
- [4] L.P. Souriau, T.D. Mast, D.L. Liu, A. Nachman and R.C. Waag, "A k-space method for large-scale models of wave propagation in tissue," *IEEE Trans. UFFC*, Submitted for publication, Feb. 1999.
- [5] M.B. Ghrist, B. Fornberg and T. Driscoll, "Staggered time integrators for wave equations," Preprint No. 394, Dept of Applied Math., Univ. of Colorado, Boulder, CO., 1998.
- [6] F.A. Duck, 1990, *Physical Properties of Tissue*, Academic Press.
- [7] L.M. Hinkelman, T.D. Mast, M.J. Orr, and R.C. Waag, "Effects of abdominal wall morphology on ultrasonic pulse distortion," *Proc. IEEE Ultrason. Symp.*, 1493-1496, 1997.
- [8] T.D. Mast, L.M. Hinkelman, M.J. Orr, V.W. Sparrow and R.C. Waag, "Simulation of ultrasonic pulse propagation through the abdominal wall," *J. Acoust. Soc. Am.*, 102, 1177-1190, 1997.
- [9] R. LaVeque, C. Zhang, "The Immersed Interface Method for Wave Equations with Discontinuous Coefficients," Univ. of Wash., Seattle WA, 1994.
- [10] F. Muir, J. Dellinger, J. Etgen and D. Nichols, "Modeling elastic fields across irregular boundaries," *Geophysics*, Vol. 57, No. 9, 1992.