



LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

Serpentine: Finite Difference Methods for Wave Propagation in Second Order Formulation

N. A. Petersson, B. Sjogreen

March 28, 2012

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

SERPENTINE: Finite Difference Methods for Wave Propagation in Second Order Formulation*

N. Anders Petersson (PI) and Bjorn Sjogreen
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory

March 26, 2012

1 Background and significance

Wave propagation phenomena are important in many DOE applications such as nuclear explosion monitoring, geophysical exploration, estimating ground motion hazards and damage due to earthquakes, non-destructive testing, underground facilities detection, and acoustic noise propagation. There are also future applications that would benefit from simulating wave propagation, such as geothermal energy applications and monitoring sites for carbon storage via seismic reflection techniques.

In acoustics and seismology, it is of great interest to increase the frequency bandwidth in simulations. In seismic exploration, greater frequency resolution enables shorter wave lengths to be included in the simulations, allowing for better resolution in the seismic imaging. In nuclear explosion monitoring, higher frequency seismic waves are essential for accurate discrimination between explosions and earthquakes. When simulating earthquake induced motion of large structures, such as nuclear power plants or dams, increased frequency resolution is essential for realistic damage predictions. Another example is simulations of micro-seismic activity near geothermal energy plants. Here, hydro-fracturing induces many small earthquakes and the time scale of each event is proportional to the square root of the moment magnitude. As a result, the motion is dominated by higher frequencies for smaller seismic events.

The above wave propagation problems are all governed by systems of hyperbolic partial differential equations in second order differential form, i.e., they contain second order partial derivatives of the dependent variables. Our general research theme in this project has been to develop numerical methods that directly discretize the wave equations in second order differential form. The obvious advantage of working with hyperbolic systems in second order differential form, as opposed to rewriting them as first order hyperbolic systems, is that the number of differential equations in the second order system is significantly smaller. Another issue with re-writing a second order system into first order form is that compatibility conditions often must be imposed on the first order form. These (Saint-Venant) conditions ensure that the solution of the first order system also satisfies the original second order system. However, such conditions can be difficult to enforce on the discretized equations, without introducing additional modeling errors.

This project has previously developed robust and memory efficient algorithms for wave propagation including effects of curved boundaries, heterogeneous isotropic, and viscoelastic materials. Partially supported by internal funding from Lawrence Livermore National Laboratory, many of these methods have been implemented in the open source software WPP [26], which is geared towards 3-D seismic wave propagation applications. This code has shown excellent scaling on up to 32,768 processors and has enabled seismic wave calculations with up to 26 Billion grid points. The WPP calculations have resulted in several publications in the field of computational seismology, e.g., [28, 2, 1, 29].

All of our current methods are second order accurate in both space and time. The benefits of higher order accurate schemes for wave propagation have been known for a long time [17, 12], but have mostly been developed for first order hyperbolic systems. For second order hyperbolic systems, it has not been known how to make finite difference schemes stable with free surface boundary conditions, heterogeneous material properties, and curvilinear coordinates.

*This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. This is Lawrence Livermore National Laboratory contribution LLNL-xyyyzz

The importance of higher order accurate methods is not necessarily to make the numerical solution more accurate, but to reduce the computational cost for obtaining a solution within an acceptable error tolerance. This is because the accuracy in the solution can always be improved by reducing the grid size h . However, in practice, the available computational resources might not be large enough to solve the problem with a low order method.

The resolution in wave propagation problems is often measured in terms of the number of grid points per shortest wave length, i.e.,

$$P_{pw} = \frac{L}{h},$$

where h is the grid size and L is the shortest wave length. In linear acoustic and elastic wave propagation, the shortest wave length and the highest significant frequency, f , are related through $L = c/f$, where c is the slowest wave speed in the material. For example, consider a case where the second order accurate method needs twice the number of grid points per wave length compared to a fourth order method. (This is a conservative estimate. There are cases where the second order method needs five times more grid points per wave length, or more.) When the fourth order method is used, the grid size can therefore be made twice as large in each spatial direction, leading to a factor of eight reduction in the number of grid points in a 3-D domain. Since the time step is proportional to the spatial grid size, it can also be made twice as large, implying that only half the number of time steps need to be taken. The fourth order stencil is about twice as wide as the second order stencil, implying that the spatial operator will be about twice as costly to evaluate per grid point. As a result, the fourth order method would in this case produce an equally accurate calculation about eight times faster, using about eight times less memory.

Even though computational resources continue to grow, they are unfortunately not unlimited. It is therefore interesting to consider how a higher order method could improve the frequency resolution on a given computer system. A fixed amount of memory translates to a maximum number of grid points in the computational grid, corresponding to a smallest grid size, h , that can be used in the calculation. Using the above relations give $P_{pw} = c/(fh)$, i.e., $f = c/(P_{pw}h)$. Hence, in the above example where a fourth order method requires half the number of grid points per wave length compared to a second order method, we get $P_{pw}^{(4)} = 0.5P_{pw}^{(2)}$. Using the same total number of grid points, the frequency resolution can therefore be doubled when the fourth order method is used.

The benefits of using a higher order accurate method increase with increasing demands on the accuracy in the numerical solution. For example, consider some wave propagation problem that needs to be solved to within error tolerance ε . To obtain this accuracy, assume that the 2nd order and a 4th order accurate method need 10 and 5 grid points per wave length, respectively. Following the above example, the 4th order method can calculate this solution about 8 times faster, using 8 times less memory, compared to the 2nd order accurate method. Now, if the error tolerance was reduced to $\varepsilon/16$, the grid size in the fourth order method would have to be reduced by a factor of two, corresponding to 10 grid points per wave length. To get the same error with the second order method, the grid size must be reduced by a factor of four, resulting in 40 grid point per wave length. For a 3-D problem, the 2nd order method therefore needs $4^3 = 64$ times more grid points and 4 times more time steps to obtain the more accurate solution. This corresponds to 64 times more memory and a factor of 256 more CPU time. In contrast, the 4th order method only needs 8 times more memory and 16 times more CPU time. Hence, since the 4th order method was already 8 times faster at computing the solution with error tolerance ε , it is 128 times faster, and needs 64 times less memory, compared with the 2nd order method, for computing the solution with error tolerance $\varepsilon/16$.

2 Summary of recent activities

To illustrate our approach for solving wave propagation problems, we show a two-dimensional cross section of a typical computational grid in Figure 1. The top curvilinear grid is made to follow the surface topography. To improve the efficiency of the method, the curvilinear grid blends smoothly into a Cartesian grid below a fixed depth. The wave speed in the earth increases with depth, and the wave lengths of the seismic waves increase proportionally. We can therefore coarsen the Cartesian grid further away from the surface, while approximately preserving the number of grid points per wave length in the numerical solution. This technique significantly reduces the total number of grid points in the numerical simulation and also allows the time step to be increased. We discretize the wave equation on this composite grid using a finite difference method that is explicit in time. The spatial discretization satisfies a summation by parts (SBP) principle on both the Cartesian [21] and the curvilinear [3] grids. The SBP principle guarantees stability of the method. We have also derived energy conserving coupling conditions for mesh refinement boundaries with hanging nodes [25]. Because the coupling conditions satisfy the SBP property, the stability of the overall method is guaranteed. Furthermore, we recently extended the approach to viscoelastic materials [27], which is important in many seismic applications.




Figure 1: A Composite grid in a vertical cross-section of the earth.

The structured nature of the finite difference grid and the explicit time-stepping method allows our method to be implemented very efficiently on modern parallel machines, and makes the finite difference method a very attractive alternative to the recently developed finite element [6], spectral element [14], discontinuous Galerkin [13, 9, 35], and finite volume [10] discretizations on unstructured grids.

We are currently generalizing our techniques to fourth order of accuracy in both space and time, and we plan to develop up to eight order accurate schemes. This generalization is possible due to our recent discovery [33] of how to construct higher order SBP finite difference approximations of second derivative terms such as $\partial/\partial x(\mu\partial u/\partial x)$, where $\mu(x) > 0$ is a variable material property and u is the dependent variable. We have designed the new discretization of the second derivative such that the requirements for obtaining an energy stable scheme for the elastic wave equation are satisfied independently of how the coefficient $\mu(x)$ varies in space. This property will also allow us to treat general curvilinear coordinate mappings.

2.1 A fourth order accurate scheme for the elastic wave equation

Several numerical methods have previously been developed for solving the elastic wave equation. The staggered grid finite difference method proposed by Virieux [34], Levander [20], Graves [11], and others, has been used extensively for seismic wave simulations. This method discretizes the elastic wave equation as a first order hyperbolic system using the velocity-stress formulation. The method is fourth order accurate in space, but only second order in time. Furthermore, the discretization of the free surface boundary condition is only second order accurate, which severely limits the accuracy of surface waves in almost incompressible materials [19]. The method uses a regular Cartesian grid with constant grid spacing, which restricts the approach to flat topographies. Analysis of the stability of the method, e.g. [20], uses Fourier techniques and is limited to the periodic problem with homogeneous material properties.

We have previously developed second order accurate discretizations of the elastic and visco-elastic wave equations, where the spatial discretization satisfies a summation by parts principle [21, 3, 25]. This technique was generalized to fourth order accuracy by Sjogreen and Petersson [33]. To satisfy the summation by parts property for the second derivatives, the weights in the discretization stencil must be modified at a number of grid points near the boundary. For the second order accurate case, the weights only need to be modified at the grid line coinciding with the boundary. For the fourth order case, the weights must be modified at the first six grid lines. Further into the domain, the discretization reverts to a centered finite difference formula. The approach generalizes to sixth and eight order discretizations, but the resulting stencils must be modified at more grid lines near the boundary, before it reverts to the corresponding centered finite difference formulae.

To fully benefit from the increased order of the spatial discretization, it is necessary to also use a higher order time-discretization. Here we are developing a modified equation approach, which can be implemented in a predictor-corrector scheme to any (even) order of accuracy. The maximum stable time steps for the scalar model problem $u_{tt} = -\kappa u$ are given in Table 1. Note that κ models the largest eigenvalue of the spatial operator, so for a fixed κ , the time step can be taken about 73 % larger with the fourth order time integration method, compared with the second order method. We have made calculations with the fourth order summation by parts scheme and the fourth order accurate time integrator, and have observed that the time step can be taken about 62 % larger with the fourth order method, compared with the second order method. Hence, the largest eigenvalue of the fourth order spatial discretization is only marginally larger than in the second order scheme. We conclude that the fourth order stencil does not add any significant stiffness to the

Order of accuracy (p)	2	4	6	8	10	12
$\Delta t_p/\Delta t_2$	1.00	1.73	1.37	2.26	1.57	2.77
Evaluations per Δt	1	2	3	4	5	6

Table 1: Maximum stable time step for the modified equation method of order p for solving $u_{tt} = -\kappa u$, $\kappa > 0$. Here, Δt_2 is the maximum time step of the second order accurate scheme. Note that the 6th and 10th order schemes are less efficient than the 2nd, 4th, 8th, and 12th order schemes.

problem. Note that time stepping stiffness is a well-known issue with higher order spectral element and discontinuous Galerkin methods, due to the clustering of internal node points near the element boundaries. This stiffness tends to be more severe as the order of accuracy increases. As our investigation indicates, this stiffness does not occur for our fourth order summation by parts discretization. Further research will tell if the trend also holds for our sixth and eight order accurate discretizations.

2.2 Boundary estimates for the elastic wave equation in almost incompressible materials

The elastic wave equation in a halfplane with a free surface boundary condition has for a long time been known to satisfy an energy estimate, and is therefore a well-posed problem. The elastic energy is a semi-norm of the solution for all heterogeneous materials with positive density, shear and bulk modulus. The energy estimate bounds this semi-norm in terms of the initial data and the internal forcing. However, the energy estimate does not provide detailed insight into how the solution depends on the material parameters, or the boundary data.

It is well known that the elastic wave equation admits compressional and shear wave solutions. These solutions are body waves. For a material with unit density, the phase velocity of the compressional and shear waves are $\sqrt{\lambda + 2\mu}$ and $\sqrt{\mu}$, respectively, where λ and μ are the Lamé parameters of the material. The elastic wave equation also admits Rayleigh surface waves. These waves travel harmonically along a free surface and decay exponentially away from the surface. The phase velocity of these waves is $\tilde{\xi}_0\sqrt{\mu}$, where $\tilde{\xi}_0 < 1$. The exact value of $\tilde{\xi}_0$ depends on λ/μ , but $\tilde{\xi}_0 \approx 0.9$. Hence, the surface waves are always slightly slower than the shear waves.

In our paper [19] we are particularly interested in materials with $\mu \ll \lambda$. According to the classical theory by Kreiss and Olinger [17], an accurate numerical solution can be guaranteed if the shortest wave length is not smaller than a constant number of grid sizes, where the constant depends on the numerical method. This theory indicates that a fixed grid size should be sufficient to maintain accuracy of the surface wave as $\mu \rightarrow 0$, assuming that the wave length of the surface wave is scaled to unity. We mention in passing that the phase velocity of a wave is defined as the ratio between its wave length and period. This scaling therefore implies that the period of the surface wave is proportional to $1/\sqrt{\mu}$ as $\mu \rightarrow 0$.

Numerical calculations indicate that the grid size predicted by the classical theory is inadequate for accurately calculating surface waves and leads to large phase errors as $\mu \rightarrow 0$. We use a normal mode analysis to explain the loss of accuracy as $\mu \rightarrow 0$, which corresponds to the incompressible limit of an elastic material. The normal mode analysis allows us to estimate the solution in terms of the boundary data, and makes the dependence on the material parameters transparent. We show that the solution is *strongly boundary stable*, except in the vicinity of the generalized eigenvalues corresponding to surface waves. Here the solution is as smooth as the boundary data, i.e., only *boundary stable* (see [18] for definitions of these stability concepts).

We use the normal mode theory to perform a modified equation analysis of the discretized boundary conditions. This analysis shows how the number of grid points per surface wave length must be increased to maintain a given error level in the numerical solution when $\mu \rightarrow 0$. When the surface waves are scaled to have unit wave length, the grid size must be proportional to $\mu^{1/2}$ for the second order method, while it suffices to take $h \sim \mu^{1/4}$ for the fourth order method. These scalings are confirmed by the numerical experiments in [19], which illustrate that the fourth order method is significantly more efficient than the second order approach.

2.3 Viscoelastic modeling of dissipative materials

Dissipative mechanisms in the earth lead to anelastic attenuation of seismic waves, which for example is important in nuclear explosion monitoring applications, where seismic signals need to be propagated over many wave lengths. The anelastic attenuation is commonly modeled by describing the earth as a viscoelastic constant- Q absorption band solid, meaning that the material has a quality factor Q , which is independent of frequency. Such material behavior can be

approximated in the time-domain by superimposing n standard linear solid (SLS) mechanisms [5].

In [27], we developed a stable finite difference approximation of the three-dimensional viscoelastic wave equation with an n -SLS material model. The proposed scheme discretizes the governing equations in second order displacement formulation using $3n$ memory variables, making it significantly more memory efficient than the commonly used first order velocity-stress formulation. The discretization is a generalization of our summation-by-parts finite difference discretization of the elastic wave equation [21, 24, 25]. Our main result is a proof that the proposed discretization is energy stable, even in the case of variable material properties. The proof relies on the summation-by-parts property of the discretization. The new scheme has been implemented with grid refinement with hanging nodes on the interface. Numerical experiments verify the accuracy and stability of the new scheme. Semi-analytical solutions for a half-space problem and the LOH.3 layer over half-space problem [8] are used to demonstrate how the number of viscoelastic mechanisms and the grid resolution influence the accuracy. We found that three standard linear solid mechanisms usually are sufficient for making the modeling error smaller than the discretization error.

2.4 Mesh refinement with hanging nodes

The compressional and shear wave lengths in the earth generally increase with depth and are often a factor of ten larger below the Moho discontinuity (at about 30 km depth), than in sedimentary basins near the surface. If a uniform grid is used to simulate the ground motion, it must have a grid spacing based on the small wave lengths near the surface, which results in over-resolving the solution at depth. As a result, the number of points in a uniform grid is unnecessarily large and the time step is unnecessarily small.

In [25] we address the over-resolution-at-depth issue by generalizing our single grid finite difference scheme to work on a composite grid consisting of a set of structured rectangular grids of different spacings, as outlined in Figure 1. Previously developed summation-by-parts properties are generalized to devise a stable second order accurate coupling of the solution across mesh refinement interfaces. By using this technique in a realistic 3-D simulation, the total number of grid points could be reduced by an order of magnitude compared to a uniform grid calculation. Alternatively, by using the same total number of grid points as in the uniform grid, the frequency content in the composite grid calculation could be doubled.

The discretization of singular source terms of point force and point moment tensor type is also studied in [25]. Based on enforcing discrete moment conditions that mimic properties of the Dirac distribution and its gradient, previous single grid formulas could be generalized to work in the vicinity of grid refinement interfaces. These source discretization formulas are shown to give second order accuracy in the solution, with the error being essentially independent of the distance between the source and the grid refinement boundary. Several numerical examples are given in [25] to illustrate the properties of our method.

2.5 Initial-boundary value problems for second order systems of partial differential equations

In [18], we describe a theory for second order hyperbolic systems based on Laplace and Fourier transform, with particular emphasis on boundary processes corresponding to generalized eigenvalues. Second order hyperbolic systems often describe problems where wave propagation is dominant. In bounded domains this leads to a large number of boundary phenomena, such as glancing waves and surface waves. Important applications of this work are the elastic wave equations and Maxwell's equations of electromagnetism. Attempts have previously been made to develop theories for second order system consisting of n equations by rewriting them as larger first order system. However, boundary phenomena such as glancing and surface waves correspond to generalized eigenvalues that are not handled by the existing theory for first order systems. Furthermore, the resulting first order system often consists of more than $2n$ equations. This leads to many complications, such as augmenting the first order system with side conditions to ensure that all its solutions also satisfy the original second order system.

Our theory uses pseudo-differential operators combined with mode analysis, and builds upon the theory for first order systems developed in [15, 16]. This approach has many desirable properties: 1) Once a second order system has been Laplace and Fourier transformed, it can always be written as a system of $2n$ first order pseudo-differential equations. Therefore, the theory of [15, 16] also applies here. 2) We can localize the problem, i.e., it is only necessary to study the Cauchy problem and halfplane problems with constant coefficients. 3) The class of problems we can treat is much larger than previous approaches based on integration by parts. 4) The relation between boundary conditions and boundary phenomena becomes transparent.

2.6 A 4th order accurate embedded boundary method

In [4], we develop a fourth-order accurate embedded boundary method for the scalar wave equation with Dirichlet or Neumann boundary conditions. The method is based on a compact Pade-type discretization of spatial derivatives together with a Taylor series method (modified equation) in time. A novel approach for enforcing boundary conditions is introduced which uses interior boundary points instead of exterior ghost points. This technique removes the small-cell stiffness problem for both Dirichlet and Neumann boundary conditions, is more accurate and robust than previous methods based on exterior ghost points, and guarantees that the solution is single-valued when slender bodies are treated. Numerical experiments are presented to illustrate the stability and accuracy of the method as well as its application to problems with complex geometries.

3 Key accomplishments

- Developed a theory for 2nd order hyperbolic partial differential equations, with particular emphasis on boundary phenomena corresponding to generalized eigenvalues, such as surface waves in the elastic wave equation [18].
- Derived 2nd order accurate discretizations of the elastic and visco-elastic wave equations that satisfy the summation by parts (SBP) principle, and therefore are energy stable [21, 24, 27].
- Generalized the 2nd order accurate SBP discretization to curvilinear grids [3], which allows free surface boundary conditions to be applied on a realistic topography.
- Generalized the 2nd order accurate SBP discretization to local mesh refinement, with hanging nodes on the refinement interface. Also derived a discretization of singular sources for such meshes [25].
- Implemented the 2nd order accurate method for solving the seismic wave propagation problem in the open source code WPP [26]. WPP has shown very good scaling to large numbers of processors. The BG/L machine at LLNL was used for our largest runs to date, which used 32,768 processors and 26 Billion grid points.
- Gave a keynote presentation on large scale seismic wave simulations at the 2008 SIAM annual meeting [22].
- Released WPP as open source software, freely downloadable from the LLNL website. The release includes a comprehensive user's guide [26].
- Explained why second order accurate numerical methods need extreme resolution for accurately computing surface waves in almost incompressible materials. Also showed why higher order of accuracy is much more efficient in this limit [19].
- Derived 4th, 6th and 8th order accurate SBP discretizations of second order derivatives with variable coefficient. These discretizations satisfy the necessary compatibility conditions with SBP discretizations of cross-terms, and lead to energy stable and higher order accurate methods for the elastic wave equation [33].

4 Transitions

Many of the numerical methods from our ASCR project are implemented in the WPP code [26]. WPP solves the 3-D elastic or viscoelastic wave equations for seismic applications on a locally refined grid. It is an open source code that has been downloaded by more than 250 users around the world since the first version was released in 2007. WPP is based on the finite difference method together with an explicit time stepping scheme, and uses MPI for communication on parallel machines. This approach is extremely well suited for parallel execution, and WPP has shown excellent performance on some of the world's largest supercomputers. WPP has shown very good scaling to large numbers of processors. The BG/L machine at LLNL was used for our largest runs to date, which used 32,768 processors and 26 Billion grid points.

4.1 Earthquake ground motions

Verification and validation is essential for establishing the reliability of computational simulation tools. In [1], we compared results from the WPP code and four other seismic simulation codes by calculating ground motions during two recent earthquakes as well as 39 scenario earthquakes on the Hayward fault in the San Francisco bay area. These simulations used a strongly heterogeneous material model of the Bay Area, developed by the U.S. Geological Survey. The validation

against seismographic recordings of the 2007, magnitude 4.18, Oakland earthquake and the 2007, magnitude 5.45, Alum Rock earthquake demonstrated good agreement between low-pass filtered seismographical measurements and simulated motions from the five different codes in the study. This agreement validates the predictive capability of the simulation codes, as well as the Bay Area material model. For the 39 scenario earthquakes on the Hayward fault, complex rupture mechanisms were developed to investigate the influence of rupture length (or magnitude), location of the hypocenter (or rupture directivity), and slip distribution. All codes in the study showed reasonably consistent results in terms of arrival times and maximum amplitudes of ground motions. Particularly good agreement was observed between WPP and the two finite elements codes in the study, which also accounted for realistic topography.

4.2 Explosion monitoring applications

In another study [29] we employed WPP to calculate high-resolution (8 Hz), three-dimensional, simulations of ground motions from shallow explosions in the presence of surface topography. This work focused on elastic propagation effects and shear wave generation. The influence on topography was investigated by filtering the actual surface elevation with Gaussian filters of different widths to generate a suite of topographies. In WPP, topography is specified independently of the grid size and the computational grid is generated in parallel during the setup phase of the program, which makes this type of parameter study easy to perform. Overall, topography was found to enhance energy propagating along the surface near the source, amplify surface waves, and balance the amount of vertically and horizontally polarized shear wave motion. All of these effects impact shear wave observations used for nuclear explosion monitoring. Further simulation studies could elucidate how the wavefield emerging from a topographically rough area ultimately propagates to regional and/or teleseismic distances.

4.3 Advanced source modeling

Explosion motions in the solid earth or atmosphere are governed by hydrodynamics and must include non-linear material response effects that result from high energy densities. However, the numerical calculation of the hydrodynamic response is well known to be computationally intensive due to non-linear constitutive behavior, especially when compared to solving linear (e.g. elastic and acoustic) wave equations. In order to propagate explosion generated ground motions from the non-linear near-source region to the far-field we have developed a hybrid modeling approach with one-way hydrodynamic-to-elastic coupling in three dimensions [36]. Near source motions are computed with GEODYN, an Eulerian hydrodynamics code with adaptive mesh refinement for high-energy loading of earth materials. Motions on a dense grid of points are saved, resampled and then passed to WPP, an anelastic finite difference code for seismic wave modeling. Our coupling strategy is based on the uniqueness theorem where motions are introduced into WPP as a boundary source and continue to propagate as elastic waves at much lower computational cost than with GEODYN. We have developed and verified the methodology to compute GEODYN hydrodynamic responses in either two- or three-dimensional domains, and pass these to WPP as 3D boundary motions on the faces of a cube. For the 2D case we compute the axisymmetric response with GEODYN and transform the elastic motions onto the coupling interface in 3D before introducing motions to WPP. The accuracy of the numerical calculations and the coupling strategy is demonstrated in cases with purely elastic medium as well as non-linear medium. Importantly we show that GEODYN can accurately model motions for a linear elastic medium including surface waves, which is essential to insure that near-source motions are correct. An application of our hybrid modeling approach is shown for a problem involving scattering by 3D heterogeneity. Our strategy by design is capable of incorporating complex non-linear effects near the source as well as volumetric and topographic material heterogeneity along the propagation path to receiver, making it very powerful for modeling a wide variety of effects and providing new prospects for modeling and understanding explosion generated seismic waveforms.

4.4 Public outreach

The WPP code was recently used to generate detailed ground motions for the April 18, 1906, magnitude $M_W \approx 7.9$, San Francisco earthquake, as well as for a $M_W = 7.05$ scenario earthquake on the Hayward fault. These ground motions will be part of the “Earthquake” show at the California Academy of Sciences in San Francisco. This show opens to the public in May of 2012 and will run for about one year. The WPP simulations will be part of a 30 minutes animation in the 75 foot Morrison planetarium, which receives about 600,000 visitors per year.

5 Technical progress (Oct 1, 2011 to April 1, 2012)

We are currently generalizing our energy stable second order accurate techniques for solving the 3-D elastic wave equation to fourth order of accuracy in both space and time.

In the following sections, we outline our approach for implementing the method on modern multi-core architectures (§ 5.1), present its performance on some test cases (§ 5.2), and discuss applications to adjoint wave propagation for seismic source estimation (§ 5.3).

5.1 Computational experiments on multi-core machines

The parallel WPP code has implemented many of the second order accurate methods from this project. In WPP, the explicit time stepping and the finite difference discretization is implemented using the MPI library in a distributed memory, message passing programming model.

Directly implementing our higher order methods in WPP would require a substantial coding effort, which is not the focus of our ASCR project. However, it is important to evaluate the performance of the numerical kernel in our high order discretizations on modern multi-core architectures. Previous arguments for advocating higher order methods have been based on the number of floating point operations per grid point [17, 12]. For modern parallel multi-core machines, memory accesses and the amount of communication are more likely to be the limiting factors.

Future supercomputers are forecasted to have a large numbers of cores per processors. To make our higher order methods perform well on multi-core processors, we will build on our success with WPP and use a hybrid programming model. This approach combines message passing between the processors with multi-threading within each multi-core processor. All calculations in an explicit finite difference solver are of local character. Multi-threading should therefore be efficient within each processor, because all cores on a processor share the same memory. However, as we shall see below, there are details in the layout of the Cache memory that must be taken into account.

To begin evaluating the performance on multi-core machines, we consider solving the elastic wave equation with our second order accurate algorithm in a two-dimensional domain. We have implemented a parallel solver in two different ways. The first implementation achieves parallelism by calls to the Pthreads library, and the second implementation uses calls to the MPI library. The computational domain is divided into patches in the same way for both implementations. Later on in this project, we will also implement the threading using the OpenMP library. The OpenMP standard provides a high level programming interface for multi-threading, where compiler directives are included in the source code on the loop level. Due to the structured layout of the finite difference grid, we expect the OpenMP instrumentation to be straight forward.

Intel Xeon X5660: The Sierra super computer system at LLNL has 1,944 computational nodes. Each node has two Intel Xeon X5660 processors, and each processor contains 6 cores, giving a total of 12 cores per node. Each core has one L1 cache memory of size 32KB and one L2 cache memory of size 256KB. There is also an L3 cache of size 12MB that is shared between the 6 cores of each processor.

Figure 2a shows the speedup on a single node, parallelized with up to 12 Pthreads, for different problem sizes. Each thread was pinned to a different core. Figure 2b shows the speedup on a single node with the same computations, but using MPI instead of Pthreads. The problem sizes ranges from 40×40 grid points up to 800×800 grid points. Noticeable in both computations is the reduced speedup for the largest problem size, displayed in cyan color in Figure 2a and 2b. A rough estimate of the memory usage showed that the 400×400 problem is small enough to fit into the L3 cache memory, while the 800×800 problem will not fit into the L3 cache. This is the likely explanation for the reduced speedup when the problem size is large. We also note that the Pthreads code is less efficient than the MPI code for the grids with 40×40 and 100×100 points. This is likely due to synchronization overhead after each time step. For the smaller problem sizes, the amount of arithmetic work is small and will not dominate the execution time. It should be possible to reduce the overhead in the Pthreads code to make its efficiency closer to the MPI code for the smaller problem sizes. However, the smaller problem sizes are not representative of the type of computations we are interested in. A realistic large scale seismic computation would be similar the 800×800 case, since a large amount of memory certainly would be needed. The current investigation indicates that computational speed would then be completely bounded by the cost of memory access. Future work on achieving optimal computational speed will therefore be focused on speeding up and reducing the amount of memory accesses.

Intel Core i7: Thread parallelism on a single Intel Core i7 processor was evaluated on the same test program as above, i.e., solving the elastic wave equation in two space dimensions using a second order accurate finite difference method

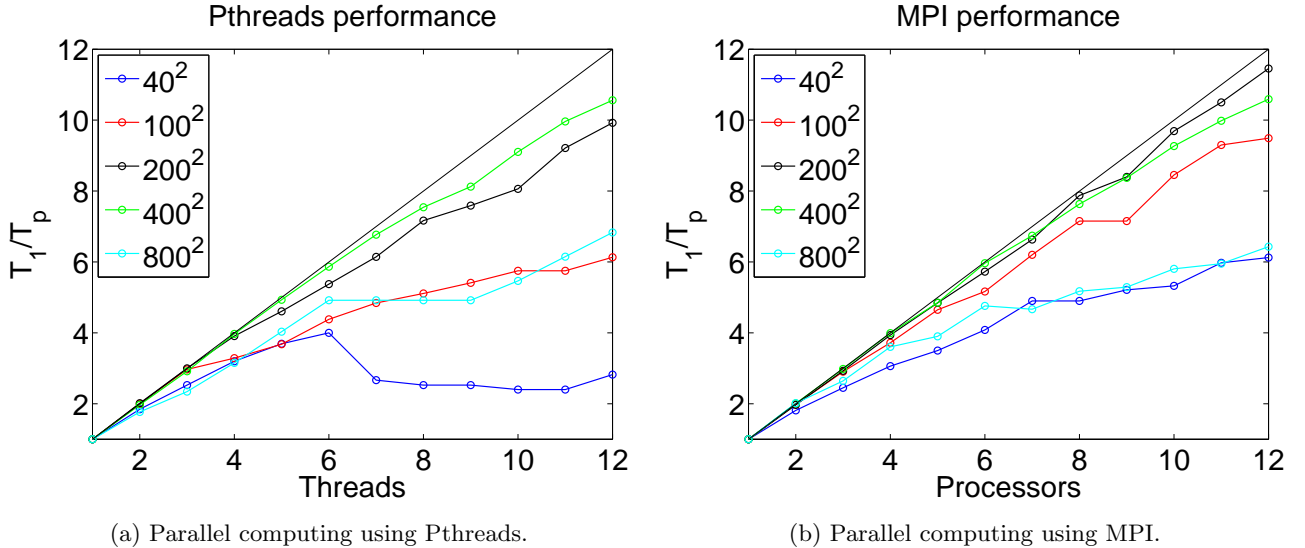


Figure 2: Speedup for two dimensional elastic wave propagation with grid sizes 40×40 (blue), 100×100 (red), 200×200 (black), 400×400 (green), and 800×800 (cyan). The solid line indicates perfect speedup.

with explicit time stepping. In this case, we fixed the grid to include 100×100 points. The Intel Core i7 chip has four cores and each core has two SMT hardware threads (not to be confused with software threads). In total there are eight logical processors. Each core has its own L1 cache memory, but the two SMT hardware threads within each core share the same L1 cache. Figure 3 shows execution time with standard partitioning of the computational domain into patches of equal size. Here we obtained a factor of three speedup on four (blue curve) and eight (red and blue curves) Pthreads. As is evident from these results, launching two Pthreads on the same core does not lead to any speedup, probably due to memory bandwidth limitations when both threads use the same L1 cache.

5.2 Comparing the second and fourth order schemes on the LOH.1 test problem


The relevant metric for comparing the efficiency of high and low order methods is the total computational cost for calculating a solution within a specified level of accuracy. The optimal order of accuracy is likely to depend on the required level of accuracy, the smoothness of the material properties and the solution, as well as the characteristics of the computer hardware. We have chosen to evaluate the total computational cost for a number of benchmark problems. Both the second and fourth order algorithms have been implemented in a message passing programming model, based on the MPI library. In our implementation, the three-dimensional computational grid is divided into patches of approximately equal size using a two-dimensional Cartesian processor decomposition (using `MPI_Cart_create`). For the second order accurate method, we use the existing implementation in WPP. It uses a two-dimensional processor decomposition because it simplifies the load balancing for a composite grid, where mesh coarsening is used in the vertical direction, see Figure 1. To make a fair comparison with the fourth order method, the same decomposition is also used in that case.

The computational seismology community has developed test cases with semi-analytical solutions, which evaluate the elastic and viscoelastic responses to moment tensor point sources. These tests use a simple layered material model [7, 8], which make them straight forward to set up. Nevertheless, the resulting ground motions are quite complicated, see Figure 4. As a result, these test cases are non-trivial to solve accurately.

The source in the LOH.1 test case is a shearing point moment tensor with a Gaussian time function with spread $\sigma = 0.06$, corresponding to a full width at half amplitude of 0.141 seconds. We define $1/0.141 \approx 7.08$ Hz to be the highest significant frequency in the solution. This test uses semi-analytical solutions to evaluate the error in the numerical three-component seismograms at stations on the free surface. We compare results of the second and fourth order method at station number 10, which is located 10 km from the epicenter. To obtain similar max errors in both simulations, the grid size in the second order method was $h = 25$ meters (max error 0.071), while $h = 50$ meters was sufficient for the fourth order method (max error 0.088). Since the lowest shear speed in the material model is 2000 m/s, these grid sizes correspond to about $P_{pw}^{(2)} = 11.3$ and $P_{pw}^{(4)} = 5.65$ grid points per shortest significant wave length. The mesh for the second order calculation used about 982 million grid points and 3485 time steps was needed to reach $t = 9$ seconds. It was run on 40 nodes (480 cores) on `sierra.lnl.gov`. The corresponding mesh for the fourth order method had about 123

Pthreads	Blue	Red
1	Core 0, SMT 0	Core 0, SMT 0
2	Core 1, SMT 0	Core 0, SMT 1
3	Core 2, SMT 0	Core 1, SMT 0
4	Core 3, SMT 0	Core 1, SMT 1
5	Core 0, SMT 1	Core 2, SMT 0
6	Core 1, SMT 1	Core 2, SMT 1
7	Core 2, SMT 1	Core 3, SMT 0
8	Core 3, SMT 1	Core 3, SMT 1

(a) Distribution of the Pthreads over the cores and SMT hardware threads.



(b) Blue and red lines correspond to the different layouts of Pthreads, see table on the left.

Figure 3: Performance on Intel Core i7 with four cores and two SMT hardware threads per core.

million grid points and 1074 time steps were taken to simulate the same motion. Hence, the 4th order calculation only needed about one eighth of the computational resources, both in terms of memory and total CPU time. We show the results from the 4th order calculation in Figure 4. However, both methods produced very similar solutions, with good overall agreements in amplitude and phase. The remaining errors are most likely attributable to interpretations of the physical properties in the material model at the discontinuity between the top layer and the underlying half-space.

Next test case will be the 3-D Rayleigh surface wave problem. This is a good model problem because there is an exact solution that can be used to evaluate the error in the numerical solution, and the material properties can be varied. It is of particular interest to evaluate the accuracy for almost incompressible materials, i.e., when the compressional speed is much larger than the shear speed.

5.3 Source estimation and adjoint wave propagation

In geothermal applications, it is of great importance to monitor micro-seismicity due to hydro-fracturing of rock. One important purpose of the monitoring is to detect interactions between existing faults and the network of fractures due to the hydro-fracturing, which may trigger significant earthquakes. The time scale of a seismic event is approximately proportional to the cube root of the total seismic moment. As a result, the motion due to smaller events is dominated by higher frequencies, which necessitates the use of higher order accurate numerical methods.

A small seismic event can be modeled by a point moment tensor source. Each source can be described by its location, start time, half-duration, and moment tensor mechanism (11 parameters in 3-D). In the following, we want to estimate the source parameters from ground motion recordings. These recordings measure the time-dependent, three-component, motion at a number of locations (usually on the surface and in bore holes). We assume that we have an accurate model of the material properties, and that the motion is governed by the isotropic elastic wave equation.

We can quantify the agreement between the recorded and simulated ground motions by the full waveform misfit,

$$\mathcal{X} = \frac{1}{2} \sum_r \int_0^T |\mathbf{u}(\mathbf{x}_r, t) - \mathbf{d}_r(t)|^2 dt. \quad (1)$$

Here, \mathbf{x}_r is the location of recording station 'r'. The numerically simulated and measured ground motions are denoted by $\mathbf{u}(\mathbf{x}_r, t)$ and $\mathbf{d}_r(t)$, respectively. Here, $\mathbf{u}(\mathbf{x}, t)$ is governed by the elastic wave equation with free surface boundary conditions along the topography. The numerical solution at the receiver locations depends implicitly on the source parameters \mathbf{p} , with components p_j , $j = 1, 2, \dots, P$. For this reason, also \mathcal{X} depends on the same parameters.

In the two-dimensional setting, the source can be described by 7 parameters. To illustrate the properties of \mathcal{X} in parameters space, we proceed as follows. We start by placing the recording stations at the locations \mathbf{x}_r in the computational model. We then use a point moment tensor source with parameters \mathbf{p}_0 and solve the elastic wave equation.




Figure 4: Radial, transverse and vertical components of the time dependent ground motion at receiver station # 10 in the LOH.1 test problem. Black and red curves show the semi-analytical and numerical solution with the 4th order method. The green curve shows the difference, i.e., the error.

The motion at each recording station is assigned to $\mathbf{d}_r(t)$. We can then modify the source parameters and compute a new numerical solution $\mathbf{u}(\mathbf{x}, t)$. By inserting it into (1), we can evaluate $\mathcal{X}(\mathbf{p})$. Unfortunately, it is difficult to plot a function of 7 parameters. By only perturbing the location of the source and keeping all other parameters at their unperturbed values, we get the contour plot in Figure 5. Note the complex behavior of the misfit function, with many ridges and valleys as well as local minima and maxima. Several of the valleys coincide with the black lines. Here, each line corresponds to source locations at a constant distance from one receiver, and are marked by its number.

Determining the source parameters from measured ground motions can be viewed as an inverse problem, where we seek to minimize $\mathcal{X}(\mathbf{p})$ under the constraint that $\mathbf{u}(\mathbf{x}, t)$ is governed by the elastic wave equation. We have developed a pre-conditioned non-linear conjugated gradient algorithm for solving this inverse problem. In each iteration, it is necessary to evaluate both the misfit and the gradient of the misfit, $\mathcal{X}'(\mathbf{p})$. The latter can conveniently be calculated by first solving one adjoint wave equation governing the adjoint wave field $\mathbf{k}(\mathbf{x}, t)$. All components of the gradient of the misfit can then be calculated by time convolutions of the form

$$\frac{\partial \mathcal{X}}{\partial p_j} = \int_0^T \int_{\Omega} \mathbf{k} \cdot \frac{\partial \mathbf{f}}{\partial p_j} d\Omega dt,$$

where $\mathbf{f}(\mathbf{x}, t) = g(t)M\nabla\delta(\mathbf{x} - \mathbf{x}_*)$ represents the point moment tensor source, located at \mathbf{x}_* .

The adjoint wave field satisfies the elastic wave equation subject to free surface boundary conditions, with point forcings at the receiver locations. Instead of initial conditions, the adjoint field is subject to terminal conditions at $t = T$, and must be solved backwards in time. The same procedure applies for the discretized elastic wave equation. Here the summation by parts property of our discretization ensures that the discrete elastic wave equation is time-reversible. The adjoint approach can therefore also be used for computing the gradient of the discretized misfit function.

To make the non-linear conjugated gradient method converge well, it is essential to scale the problem appropriately. We therefore introduce a scaled parameter vector $\hat{\mathbf{p}} = E^T \mathbf{p}$. It is well known that the scaling should approximate the Hessian of the misfit in the sense $EE^T \approx \mathcal{X}''$. In the 2-D setting, the Hessian can be calculated by solving seven additional elastic wave equations. Our preliminary investigation indicates that a good balance between the additional cost and the accelerated convergence is obtained by calculating the Hessian for one parameter value, and using the fixed diagonal scaling $E_{jj} = \sqrt{\mathcal{X}''_{jj}}$ throughout the iteration. In a test case, the conjugated gradient algorithm converges in about 25-30 iterations. The computational effort for performing this calculation is approximately equivalent to 100 forward (in time)




Figure 5: Countour plot of the misfit \mathcal{X} as function of the source location, when all other parameters have the correct values. The minima is at the center of the dark blue area.

solutions of the elastic wave equation.

This preliminary work has been presented at a number of seminars both at LLNL and externally [23]. We have so far been solving the inverse problem in 2-D, are currently generalizing the technique to 3-D.

5.4 Collaborations

Since the beginning of this project, we have worked closely with Prof. Kreiss, who is a professor emeritus from UCLA and the Royal Institute of Technology in Stockholm, Sweden. Prof. Kreiss has helped with the development of our approach for hyperbolic problems in second order formulation and has made many important theoretical contributions to our work. His contributions will continue to be important as we approach more challenging stability and accuracy questions for more complicated systems of second order hyperbolic PDEs, including effects of anisotropy and higher order accurate methods for visco-elastic wave propagation.

We have also worked extensively with Dr. Appelo, who is an expert in far field boundary conditions for wave propagation problems. Dr. Appelo was first a postdoctoral fellow in our project, and helped develop the curvilinear technique for satisfying free surface boundary conditions on realistic topographies. Together with him, we have also developed super-grid scale far-field boundary conditions for the elastic wave equation. This class of boundary conditions are more robust and less expensive compared with the perfectly matched layers (PML), which often are used in electromagnetic wave propagation. Dr. Appelo is now a professor of applied mathematics at the University of New Mexico in Albuquerque.

During this project, we have also worked closely with several computational seismologists at LLNL, for example, Dr. Rodgers, Dr. Harben, Dr. Pitarka, Dr. Mellors, Dr. Morency, and Dr. Harris. The interaction with applied researchers has significantly improved our understanding of how wave simulations are being used in their field. These interactions have been critical in shaping the WPP code.

6 Work plan for the rest of FY12 and FY13

Compared to the original milestones in the proposal, we have prioritized our work on adjoint wave propagation. This is motivated by the geothermal application discussed above, where we estimate source parameters based on measurements of ground motion. This work is leveraged by an internally funded LDRD project at LLNL.

During the remainder of FY12, we will continue to evaluate the performance of our fourth order SBP method on benchmark problems, such as the aforementioned Rayleigh surface wave problem. We also plan to implement the 6th and 8th order generalizations of our spatial discretizations. We will continue to evaluate performance on new multi-core machines as they become available. For example the TLCC2 and Sequia super-computers, which are both scheduled to be operational by Spring/Summer of 2012.

For FY13 we will generalize our higher order technique to work on a composite grid, see Figure 1. For the second order accurate method, the mesh refinement near the free surface allowed the frequency resolution to be doubled in practical calculation. If we can obtain the same advantage with our higher order SBP approach, our methods will allow the frequency content in seismic wave simulation to be made significantly higher than ever before. We will also generalize the SBP approach to work on curvilinear meshes, to allow the boundary conditions to be imposed on a realistic topography. As these more complex methods are implemented, we will keep evaluating the computational efficiencies on modern multi-core machines, and continue our investigation of a mixed multi-threaded / message passing programming model. We also plan to develop a higher order accurate approach for the viscoelastic wave equation and evaluate its performance on the LOH.3 test problem.

References

- [1] B. T. Aagard, R. Graves, A. Rodgers, T. M. Brocher, R. W. Simpson, D. Dreger, N. A. Petersson, S. C. Larsen, S. Ma, and R. C. Jachens. Ground motion modeling of Hayward fault scenario earthquakes, part II: Simulation of long-period and broadband ground motions. *Bull. Seism. Soc. Amer.*, 100(6):2945–2977, 2010.
- [2] B. Aagardh, T. Brocher, D. Dreger, A. Frankel, R. Graves, S. Harmsen, S. Larsen, K. McCandless, S. Nilsson, N. A. Petersson, A. Rodgers, B. Sjogreen, and M. L. Zoback. Ground motion modeling of the 1906 San Francisco earthquake II: Ground motion estimates for the 1906 earthquake and scenario events. *Bull. Seism. Soc. Amer.*, 98:1012–1046, 2008.
- [3] D. Appelö and N. A. Petersson. A stable finite difference method for the elastic wave equation on complex geometries with free surfaces. *Comm. Comput. Phys.*, 5:84–107, 2009.
- [4] D. Appelö and N. A. Petersson. A fourth-order accurate embedded boundary method for the wave equation. Technical report LLNL-JRNL-417163, Lawrence Livermore National Laboratory, 2012. (to appear in *SIAM J. Sci. Comput.*).
- [5] J. M. Carcione. *Wave fields in real media: wave propagation in anisotropic, anelastic and porous media*, volume 31 of *Handbook of geophysical exploration: seismic exploration*. Pergamon, Elsevier Science, 2001.
- [6] G. Cohen and S. Fauqueux. Mixed spectral finite elements for the linear elasticity system in unbounded domains. *SIAM J. Sci. Comput.*, 26(3):864–884, 2005.
- [7] S. M. Day, J. Bielak, D. Dreger, S. Larsen, R. Graves, A. Pitarka, and K. B. Olsen. Test of 3D elastodynamic codes: Lifelines project task 1A01. Technical report, Pacific Earthquake Engineering Center, 2001.
- [8] S. M. Day, J. Bielak, D. Dreger, S. Larsen, R. Graves, A. Pitarka, and K. B. Olsen. Test of 3D elastodynamic codes: Lifelines program task 1A02. Technical report, Pacific Earthquake Engineering Center, 2003.
- [9] M. Dumbser and M. Käser. An arbitrary high-order discontinuous Galerkin method for elastic waves on unstructured meshes - II. The three-dimensional isotropic case. *Geophys. J. Int.*, 167(1):319–336, 2006.
- [10] M. Dumbser, M. Käser, and J. de la Puente. Arbitrary high-order finite volume schemes for seismic wave propagation on unstructured meshes in 2D and 3D. *Geophys. J. Int.*, 171(2):665–694, 2007.
- [11] R.W. Graves. Simulating seismic-wave propagation in 3-D elastic media using staggered-grid finite differences. *Bull. Seismo. Soc. Amer.*, 86(4):1091–1106, 1996.
- [12] B. Gustafsson, H.-O. Kreiss, and J. Olinger. *Time dependent problems and difference methods*. Wiley-Interscience, 1995.
- [13] M. Käser and M. Dumbser. An arbitrary high-order discontinuous Galerkin method for elastic waves on unstructured meshes - I. The two-dimensional isotropic case with external source terms. *Geophys. J. Int.*, 166(2):855–877, 2006.
- [14] D. Komatitsch and J. Tromp. Introduction to the spectral element method for three-dimensional seismic wave propagation. *Geophys. J. Int.*, 139:806–822, 1999.
- [15] H.-O. Kreiss. Initial boundary value problems for hyperbolic systems. *Commun. Pur. Appl. Math.*, 23:277–298, 1970.
- [16] H.-O. Kreiss and J. Lorenz. *Initial-Boundary Value Problems and the Navier-Stokes Equations*. Academic Press, 1989.
- [17] H.-O. Kreiss and J. Olinger. Comparison of accurate methods for the integration of hyperbolic equations. *Tellus*, 24:199–215, 1972.
- [18] H.-O. Kreiss, O.E. Ortiz, and N.A. Petersson. Initial-boundary value problems for second order systems of partial differential equations. *Math. Model. Numer. Anal.*, 46(3):559–593, 2012.
- [19] H.-O. Kreiss and N.A. Petersson. Boundary estimates for the elastic wave equation in almost incompressible materials. LLNL-JRNL 482152, Lawrence Livermore National Laboratory, 2011. (Submitted to *SIAM J. Numer. Anal.*).

- [20] A.R. Levander. Fourth-order finite-difference P-SV seismograms. *Geophysics*, 53:1425–1436, 1988.
- [21] S. Nilsson, N. A. Petersson, B. Sjögreen, and H.-O. Kreiss. Stable difference approximations for the elastic wave equation in second order formulation. *SIAM J. Numer. Anal.*, 45:1902–1936, 2007.
- [22] N. A. Petersson. Large scale simulations of earthquakes. Invited keynote presentation, SIAM Annual Meeting, San Diego, CA, July 2008. LLNL-PRES-405138.
- [23] N. A. Petersson. Source estimation by full wave form inversion. Invited seminars at LLNL, LBNL, USC, and SDSU, Jan-Apr 2012. LLNL-PRES-525651.
- [24] N. A. Petersson and B. Sjögreen. An energy absorbing far-field boundary condition for the elastic wave equation. *Comm. Comput. Phys.*, 6:483–508, 2009.
- [25] N. A. Petersson and B. Sjögreen. Stable grid refinement and singular source discretization for seismic wave simulations. *Comm. Comput. Phys.*, 8(5):1074–1110, November 2010.
- [26] N. A. Petersson and B. Sjögreen. User’s guide to WPP version 2.1. Technical Report LLNL-SM-487431, Lawrence Livermore National Laboratory, 2011. (Source code available from <https://computation.llnl.gov/casc/serpentine>).
- [27] N. A. Petersson and B. Sjögreen. Stable and efficient modeling of anelastic attenuation in seismic wave propagation. *Comm. Comput. Phys.*, 12(1):193–225, 2012.
- [28] A. Rodgers, N. A. Petersson, S. Nilsson, B. Sjögreen, and K. McCandless. Broadband waveform modeling of moderate earthquakes in the San Francisco bay area and preliminary assesment of the USGS 3D seismic velocity model. *Bull. Seism. Soc. Amer.*, 98:969–988, 2008.
- [29] A. Rodgers, N. A. Petersson, and B. Sjögreen. Simulation of topographic effects on seismic waves from shallow explosions near the North Korean nuclear test site with emphasis on shear wave generation. *J. Geophys. Res. Sol. Earth*, 115:B11309, 2010.
- [30] A. J. Rodgers, E. Matzel, M. Pasyanos, N. A. Petersson, B. Sjögreen, C. Bono, O. Vorobiev, T. Antoun, and W. Walter. Seismic simulations using parallel computing and three-dimensional earth models to improve nuclear explosion phenomenology and monitoring. Portsmouth, VA, September 2008. Monitoring Research Review.
- [31] A. J. Rodgers, O. Y. Vorobiev, N. A. Petersson, B. A. Sjögreen, and M. Avants. Toward end-to-end seismic simulations for nuclear explosion monitoring. Tucson, AZ, September 2009. Monitoring Research Review.
- [32] A. J. Rodgers, O. Y. Vorobiev, B. Sjögreen, and N. A. Petersson. Simulation of seismic waves from underground explosions in geologic media: FY2009 progress report. Technical report LLNL-TR-420208, Lawrence Livermore National Laboratory, 2009.
- [33] B. Sjögreen and N. A. Petersson. A fourth order accurate finite difference scheme for the elastic wave equation in second order formulation. LLNL-JRNL 483427, Lawrence Livermore National Laboratory, 2011. (to appear in *J. Sci. Comput.*).
- [34] J. Virieux. P-SV wave propagation in heterogeneous media: Velocity-stress finite-difference method. *Geophysics*, 51:889–901, 1986.
- [35] L.C. Wilcox, G. Stadler, C. Burstedde, and O. Ghattas. A high-order discontinuous Galerkin method for wave propagation through coupled elastic-acoustic media. *J. Comput. Phys.*, 229:9373–9396, 2010.
- [36] H. Xu, A.J. Rodgers, I.N. Lomov, N.A. Petersson, B. Sjögreen, and O.Y. Vorobiev. Simulation of explosion ground motion using a hydrodynamic-to-elastic coupling approach in three dimensions. Technical Report LLNL-JRNL-xxxxxyy, Lawrence Livermore National Laboratory, 2012. (submitted to *Bull. Seismo. Soc. Amer.*).