

Parallel Octree-Based Finite Element Method for Large-Scale Earthquake Ground Motion Simulation

J. Bielak¹, O. Ghattas², and E.-J. Kim³

Abstract: We present a parallel octree-based finite element method for large-scale earthquake ground motion simulation in realistic basins. The octree representation combines the low memory per node and good cache performance of finite difference methods with the spatial adaptivity to local seismic wavelengths characteristic of unstructured finite element methods. Several tests are provided to verify the numerical performance of the method against Green's function solutions for homogeneous and piecewise homogeneous media, both with and without anelastic attenuation. A comparison is also provided against a finite difference code and an unstructured tetrahedral finite element code for a simulation of the 1994 Northridge Earthquake. The numerical tests all show very good agreement with analytical solutions and other codes. Finally, performance evaluation indicates excellent single-processor performance and parallel scalability over a range of 1 to 2048 processors for Northridge simulations with up to 300 million degrees of freedom.

keyword: Earthquake ground motion modeling, octree, parallel computing, finite element method, elastic wave propagation

1 Introduction

Wave propagation simulations for earthquake-induced ground motion have been performed for over 30 years to gain a better understanding of the distribution of the earthquake motion in urban regions in space and time. Such insight has contributed to the development of building codes in which a seismic-prone region is divided into different zones of comparable seismic hazard, with the goal of reducing seismic risk. The dramatic improvement in supercomputing performance has more recently enabled seismologists and earthquake engineers to more

accurately understand the effects of source, wave propagation, and local site conditions on the ground motion. In particular, using parallel computers with several thousand processors, it has now become possible to model ground motion in large, highly heterogeneous basins, such as the Los Angeles (LA) basin, with sufficient resolution to capture frequencies of interest, for realistic geological models.

An earthquake ground motion simulation entails solving numerically the elastodynamic wave equations. There are several numerical methods available for ground motion simulations. The finite difference method (FDM), the boundary element method (BEM) and the finite element method (FEM) are commonly used. In seismology and earthquake engineering, the FDM has been the most popular technique due to its satisfactory accuracy, ease of implementation, and low memory needed per grid point [e.g., Virieux (1984); Levander (1988); Graves (1996); Pitarka, Irikura, Iwata, and Sekiguchi (1998)]. A number of earthquake ground motion simulations in the greater LA basin have been computed using the FDM [e.g., Vidale and Helmberger (1988); Frankel and Vidale (1992); Yomogida and Etgen (1993); Schrivner and Helmberger (1994); Olsen and Archuleta (1996); Graves (1998)]. However, for heterogeneous media with large contrasts in material stiffness, the (conventional form of the) FDM suffers due to its reliance on a regular grid, as illustrated in Figure 1a. The regular grid is not capable of adapting to the local wavelengths of propagating waves, which are shorter for softer materials. Thus, grid resolution is dictated by the shortest wavelengths, which results in over-refined grids in stiffer regions. In addition, material interfaces are resolved with $O(h)$ geometric error, unless the interfaces are axis-aligned. Finally, the time step in an explicit time integrator (which is most commonly used in wave propagation) is artificially small to accommodate the CFL stability condition in the over-refined stiff regions.

The main advantage of the BEM is its unique ability

¹ Carnegie Mellon University, Pittsburgh, PA, USA.

² University of Texas at Austin, Austin, TX, USA.

³ Duke University, Durham, NC, USA.




Figure 1 : Examples of spatial discretization of a basin for different numerical methods. (a) Regular grid [Pitarka, Irikura, Iwata, and Sekiguchi (1998)]. (b) Unstructured tetrahedral mesh [Bao, Bielak, Ghattas, Kallivokas, O'Hallaron, Schewchuk, and Xu (1998)]. (c) Octree mesh [Kim, Bielak, and Ghattas (2003)].

to provide a complete solution in terms of boundary values only, with substantial savings in modeling effort. Zang and Chopra (1991) used a BEM approach to study the 3D response of a canyon in an elastic half-space. Others have investigated different types of integral equations and their numerical treatment [e.g., Sánchez-Sesma (1978); Wong (1979), Sánchez-Sesma and Esquivel (1979); Wong (1982); Dravinski (1982); Sánchez-Sesma, Bravo, and Herrera (1985); Kawase (1988)]. But the BEM is not appropriate for ground motion simulations in heterogeneous basins, since it requires piecewise-homogeneity of the domain.

On the other hand, the FEM is designed to be effective for heterogeneous media. Unlike the FDM, the FEM is capable of adapting the mesh to local features of the solution, as shown in Figure 1b. Thus over-refinement is avoided and longer time steps can be taken, since the CFL stability limit is on the order of that required for accuracy. These advantages become more dramatic with increasing contrast in shear wave velocity in the heterogeneous medium. However, the FEM comes with several disadvantages relative to FDM: gen-

erating an appropriate wavelength-adaptive mesh can require considerable effort; storing the system stiffness matrix requires considerable memory, since the “stencil” varies from node to node; and indirect addressing associated with unstructured meshes results in generally poorer cache performance. In fact, these disadvantages are not intrinsic to the FEM; they are associated with unstructured mesh methods in general. For regular grids, the FEM can be implemented with all of the advantages of the FDM. Because of the difficulties in meshing complex heterogeneous basins, there are fewer examples of FEM-based earthquake ground motion simulations [e.g., Toshinawa and Ohmachi (1988); Li, Bielak, and Ghattas (1992); Bao, Bielak, Ghattas, Kallivokas, O'Hallaron, Shewchuk, and Xu (1996); Bao, Bielak, Ghattas, Kallivokas, O'Hallaron, Schewchuk, and Xu (1998); Komatitsch and Tromp (1999); Yoshimura, Bielak, Hisada, and Fernandez (2003); Akcelik, Bielak, Biros, Epanomeritakis, Fernandez, Ghattas, Kim, Lopez, O'Hallaron, Tu, and Urbanic (2003)].

In this article, we present, verify, and demonstrate an octree-based finite element method for seismic wave propagation that overcomes the difficulties of the unstructured mesh-based FEM cited above [Kim (2003)]. The octree method combines the adaptivity of the FEM to local wavelengths with the low memory requirements and good cache performance associated with the FDM. The keys are an octree data structure for generating, representing, and refining the mesh (as illustrated in Figure 1c); a matrix-free implementation enabled by the regular shape of the elements; and enforcement of algebraic constraints at hanging nodes to maintain a conforming approximation across element interfaces. Section 2 describes the ingredients of the octree method, while Section 3 discusses the discretization and solution of the elastic wave propagation equation. Incorporation of anelastic attenuation is described in Section 4. Section 5 presents numerical verification tests of the octree method, while Section 6 provides performance and scalability results to 2048 processors.

2 Octree-based finite element method

Octrees have been used as a basis for finite element approximation since at least the early 90s [Young, Melvin, Bieterman, Johnson, Samant, and Bussolletti (1991)]. Our interest in octrees stems from their ability to adapt to the wavelengths of propagating seismic waves while

maintaining a regular shape of finite elements. here, leaves associated with the lowest level of the octree are identified with trilinear hexahedral finite elements and used for a Galerkin approximation of a suitable weak form of the elastic wave propagation equation. The hexahedra are recursively subdivided into 8 elements until a local refinement criterion is satisfied. For seismic wave propagation in heterogeneous media, the criterion is that the longest element edge should be such that there result at least p nodes per local shear wavelength, as determined by the local shear wave velocity β and the maximum frequency of interest f_{max} . In other words, $h_{max} < \frac{\beta}{pf_{max}}$. For trilinear hexahedra and taking into account the accuracy with which we know typical basin properties, we take typically $p = 10$. An additional condition that drives mesh refinement is that the element size not differ by more than a factor of two across neighboring elements (the octree is then said to be *balanced*). Note that as in finite difference methods, the octree does not explicitly represent material interfaces within the earth, and instead accepts $O(h)$ error in representing them implicitly. This is usually justified for earthquake modeling since the location of interfaces is known at best to the order of the seismic wavelength. If needed, higher-order accuracy in representing arbitrary interfaces can be achieved by local adjustment of the finite element basis [e.g., Young, Melvin, Bieterman, Johnson, Samant, and Bussoletti (1991)].

Figure 2 depicts the octree mesh (and its 2D counterpart, a quadtree). The left drawing illustrates a factor-of-two edge length difference (a “legal” refinement) and a factor-of-four difference (an “illegal” refinement). Unless additional measures are taken, so-called *hanging nodes* that separate different levels of refinement (indicated by solid circles and the subscript h in the figure) result in a possibly discontinuous field approximation, which can destroy the convergence properties of the Galerkin method. Several possibilities exist to remedy this situation by enforcing continuity of displacement field across the interface either strongly (e.g., by construction of special transition elements) or weakly (e.g., via mortar elements or discontinuous Galerkin approximation). The simplest technique is to enforce continuity by algebraic constraints that require the displacement of the hanging node be the average of the displacement of its *anchored* neighbors (indicated by open circles and the subscript a). As illustrated in Figure 2, the displacement




Figure 2 : Quadtree and Octree-Based Mesh

of an edge hanging node, $\bar{u}_h$, should be the average of its two edge neighbors u_a^i and u_a^j , and the displacement of a face hanging node, $\hat{u}_h$, should be the average of its four face neighbors u_a^i , u_a^j , u_a^k , and u_a^l . Efficient implementation of these algebraic constraints will be discussed in the next section. As evident from the figure, when the octree is balanced, an anchored node cannot also be a hanging node.

Our earlier earthquake modeling code was based on an unstructured mesh data structure and linear tetrahedral finite elements [Bao, Bielak, Ghattas, Kallivokas, O’Hallaron, Shewchuk, and Xu (1996); Bao, Bielak, Ghattas, Kallivokas, O’Hallaron, Schewchuk, and Xu (1998)]. Compared to that approach, the present octree-based method has several important advantages:

- The octree meshes are much more easily generated than general unstructured tetrahedral meshes, particularly when the number of elements increases above 50 million. We use an efficient out-of-core octree-based hexahedral mesh generator [Tu, O’Hallaron, and López (2002); Tu and O’Hallaron (2004)] that can generate meshes of sizes that are limited only by available disk space. (Since each basin is meshed just once for a given resolution of interest—but subjected to many earthquake scenarios—mesh generation can be done off-line.)
- The hexahedra provide somewhat greater accuracy per node (the asymptotic convergence rate is unchanged, but the constant is typically improved over tetrahedral approximation).
- The hexahedra all have the same form of the element stiffness matrices, scaled simply by element

size and material properties (which are stored as vectors), and thus no matrix storage is required at all. This results in a substantial decrease in required memory—about an order of magnitude, compared to our node-based tetrahedral code.

- Because of the matrix-free implementation, (stiffness) matrix-vector products are carried out at the element level. This produces much better cache utilization by relegating the work that requires indirect addressing (and is memory bandwidth-limited) to vector operations, and recasting the majority of the work of the matrix-vector product as local element-wise dense matrix computations. The result is a significant boost in performance.

These features permit earthquake simulations to substantially higher frequencies and lower resolved shear wave velocities than heretofore possible. In the next section, we describe the octree-based discretization and solution of the elastic wave equation.

3 Wave propagation model and octree discretization

We model seismic wave propagation in the earth via Navier's equations of linear elastodynamics. Let $\mathbf{u}$ represent the vector field of the three displacement components, λ and μ the Lamé moduli and ρ the density distribution, $\mathbf{b}$ a time-dependent body force representing the seismic source, and $\mathbf{t} = [\mu(\nabla\mathbf{u} + \nabla\mathbf{u}^T) + \lambda(\nabla \cdot \mathbf{u})\mathbf{I}] \mathbf{n}$ the surface traction vector. Let Ω be an open bounded domain in $\mathbb{R}^3$ with free surface Γ_{FS} , truncation boundary Γ_{AB} , and outward unit normal to the boundary $\mathbf{n}$. The initial-boundary value problem is then written as:

$$\begin{aligned} \rho \ddot{\mathbf{u}} - \nabla \cdot [\mu(\nabla\mathbf{u} + \nabla\mathbf{u}^T) + \lambda(\nabla \cdot \mathbf{u})\mathbf{I}] &= \mathbf{b} \\ \text{in } \Omega \times (0, T), \\ \mathbf{n} \times \mathbf{n} \times \mathbf{t} &= \mathbf{n} \times \mathbf{n} \times \dot{\mathbf{u}} \sqrt{\rho\mu} \quad \text{on } \Gamma_{AB} \times (0, T), \\ \mathbf{n} \cdot \mathbf{t} &= \mathbf{n} \cdot \dot{\mathbf{u}} \sqrt{\rho(\lambda + 2\mu)} \quad \text{on } \Gamma_{AB} \times (0, T), \\ \mathbf{t} &= \mathbf{0} \quad \text{on } \Gamma_{FS} \times (0, T), \\ \mathbf{u} &= \mathbf{0} \quad \text{in } \Omega \times \{t = 0\}, \\ \dot{\mathbf{u}} &= \mathbf{0} \quad \text{in } \Omega \times \{t = 0\}, \end{aligned} \quad (1)$$

With this model, p waves propagate with velocity $\alpha = \sqrt{(\lambda + 2\mu)/\rho}$, and s waves with velocity $\beta = \sqrt{\mu/\rho}$. The continuous form above does not include material attenuation, which we introduce at the discrete level via a

Rayleigh damping model, as discussed below. The vector $\mathbf{b}$ comprises a set of body forces that equilibrate an induced displacement dislocation on a fault plane, providing an effective representation of earthquake rupture on the plane. For example, for a seismic excitation idealized as a point source, $\mathbf{b} = -\mu\nu\mathbf{A}\mathbf{M}f(t)\nabla\delta(\mathbf{x} - \boldsymbol{\xi})$ [Aki and Richards (1980)]. In this expression, ν is the average earthquake dislocation; A the rupture area; $\mathbf{M}$ the (normalized) seismic moment tensor, which depends on the orientation of the fault; $f(t)$ the (normalized) time evolution of the rupture; and $\boldsymbol{\xi}$ the source location.

Since we model earthquakes within a portion of the earth, we require appropriately positioned absorbing boundaries to account for the truncated exterior. For simplicity, in (1) the absorbing boundaries are given as dashpots on Γ_{AB} , which approximate the tangential and normal components of the surface traction vector $\mathbf{t}$ with time derivatives of corresponding components of the displacement vector. Even though this absorbing boundary is approximate, it is local in both space and time, which is particularly important for large-scale parallel implementation. Finally, we enforce traction-free conditions on the free surface of the earth.

We apply standard Galerkin finite element approximation in space to the appropriate weak form of the initial-boundary value problem (1). Let $\mathcal{U}$ be the space of admissible solutions (which depends on the regularity of $\mathbf{b}$), $\mathcal{U}_h$ be a finite element subspace of $\mathcal{U}$, and $\mathbf{v}_h$ be a test function from that subspace. Then the weak form is written as follows.

Find $\mathbf{u}_h \in \mathcal{U}_h$ such that

$$\begin{aligned} \int_{\Omega} [\rho \ddot{\mathbf{u}}_h \cdot \mathbf{v}_h + \frac{\mu}{2} (\nabla\mathbf{u}_h + \nabla\mathbf{u}_h^T) \cdot (\nabla\mathbf{v}_h + \nabla\mathbf{v}_h^T)] dx \\ + \int_{\Omega} [\lambda(\nabla \cdot \mathbf{u}_h)(\nabla \cdot \mathbf{v}_h) - \mathbf{b} \cdot \mathbf{v}_h] dx \\ = \int_{\Gamma_{AB}} \sqrt{\rho\mu} (\mathbf{n} \times \mathbf{n} \times \dot{\mathbf{u}}_h) \cdot (\mathbf{n} \times \mathbf{n} \times \mathbf{v}_h) ds \\ + \int_{\Gamma_{AB}} \sqrt{\rho(\lambda + 2\mu)} (\mathbf{n} \cdot \dot{\mathbf{u}}_h) (\mathbf{n} \cdot \mathbf{v}_h) ds, \quad \forall \mathbf{v}_h \in \mathcal{U}_h. \end{aligned} \quad (2)$$

Finite element approximation is effected via piecewise trilinear basis functions and associated trilinear hexahedral elements on an octree mesh. This strikes a balance between simplicity, low memory (since all element stiffness matrices are the same modulo scale factors), and reasonable accuracy.⁴

⁴ The output quantities of greatest interest are displacements and

Upon spatial discretization, we obtain a system of ordinary differential equations of the form

$$\begin{aligned} \mathbf{M}\ddot{\mathbf{u}} + (\mathbf{C}^{AB} + \mathbf{C}^{att})\dot{\mathbf{u}} + \mathbf{K}\mathbf{u} &= \mathbf{b}, \\ \dot{\mathbf{u}}(0) &= \mathbf{0}, \\ \mathbf{u}(0) &= \mathbf{0}, \end{aligned} \quad (3)$$

where $\mathbf{M}$ and $\mathbf{K}$ are mass and stiffness matrices, arising from the terms involving ρ and (μ, λ) in (2), respectively; $\mathbf{b}$ is a body force vector resulting from a discretization of the seismic source model; and damping matrix $\mathbf{C}^{AB}$ reflects contributions of the absorbing boundaries. We have also introduced the damping matrix $\mathbf{C}^{att}$ to simulate the effect of energy dissipation due to anelastic material behavior; it consists of a linear combination of mass and stiffness matrices and its form will be discussed in the next section.

The time dimension is discretized using central differences. The algorithm is made explicit using a diagonalization scheme that (1) lumps the mass matrix $\mathbf{M}$, the absorbing boundary matrix $\mathbf{C}^{AB}$, and the mass component of the the material attenuation matrix $\mathbf{C}^{att}$ (all of which have Gram structure and are therefore spectrally equivalent to the identity), and (2) splits the diagonal and off-diagonal portions of the stiffness component of $\mathbf{C}^{att}$, time-lagging the latter. The resulting update for the displacement field at time step $k + 1$ is given by

$$\begin{aligned} \left[\mathbf{M} + \frac{\Delta t}{2} \mathbf{C}_{diag}^{AB} + \frac{\Delta t}{2} \mathbf{C}_{diag}^{att} \right] \mathbf{u}_{k+1} \\ = \left[2\mathbf{M} - \Delta t^2 \mathbf{K} - \frac{\Delta t}{2} \mathbf{C}_{off}^{AB} - \frac{\Delta t}{2} \mathbf{C}_{off}^{att} \right] \mathbf{u}_k \\ + \left[\frac{\Delta t}{2} \mathbf{C}^{AB} + \frac{\Delta t}{2} \mathbf{C}^{att} - \mathbf{M} \right] \mathbf{u}_{k-1} + \Delta t^2 \mathbf{b}_k. \end{aligned} \quad (4)$$

The time increment Δt must satisfy a local CFL condition for stability. Space is discretized over the octree mesh (each leaf corresponds to a hexahedral element) that resolves local seismic wavelengths as discussed above. This insures that the CFL-limited time step is of the order of that needed for accuracy, and that excessive dispersion errors do not arise due to over-refined meshes.

Spatial discretization via refinement of an octree produces a non-conforming mesh, resulting in a discontinuous displacement approximation. We restore continuity of the displacement field across refinement interfaces by

velocities, as opposed to stresses.

imposing algebraic constraints that require the displacement at a hanging node to be consistent with the approximation along the neighboring element face or edge, as discussed above. We can express these algebraic continuity constraints in the form

$$\mathbf{u} = \mathbf{B}\tilde{\mathbf{u}},$$

where $\tilde{\mathbf{u}}$ denotes the displacements at the non-hanging (i.e., independent) nodes, and $\mathbf{B}$ is a sparse constraint matrix. In particular, $\mathbf{B}_{ij} = \frac{1}{4}$ if hanging node i is a face neighbor of anchored node j and $\frac{1}{2}$ if it is an edge neighbor; $\mathbf{B}_{ij} = 1$ simply identifies a non-hanging node; and $\mathbf{B}_{ij} = 0$ otherwise. Rewriting the linear system (4) as

$$\mathbf{A}\mathbf{u}_{k+1} = \mathbf{c}(\mathbf{u}_k, \mathbf{u}_{k-1})$$

we can impose the continuity constraints via the projection

$$\mathbf{B}^T \mathbf{A} \mathbf{B} \tilde{\mathbf{u}}_{k+1} = \mathbf{B}^T \mathbf{c}(\mathbf{u}_k, \mathbf{u}_{k-1}). \quad (5)$$

The reduced matrix $\mathbf{B}^T \mathbf{A} \mathbf{B}$ is then further lumped to render it diagonal, so that the constrained update (5) is explicit. In fact the resulting reduced matrix can be constructed simply by dividing the contributions of the hanging nodes and adding them in equal portions to diagonal components of the corresponding anchored nodes. The righthand side of (5) is determined at each time step by computing $\mathbf{c}(\mathbf{u}_k, \mathbf{u}_{k-1})$, i.e. the righthand side of (4), in an element-by-element fashion, and then applying the reduction $\mathbf{B}^T \mathbf{c}$. This amounts again to dividing the contributions of the hanging components of $\mathbf{c}$ equally among the corresponding anchored nodes. The work involved in enforcing the constraints is proportional to the number of hanging nodes, which can be a sizable fraction of the overall number of nodes for a highly irregular octree, but is at most of $O(N)$. Therefore, the per-iteration complexity of the update (5) remains linear in the number of nodes.

The combination of an octree-based wavelength-adaptive mesh, piecewise trilinear Galerkin finite elements in space, explicit central differences in time, constraints that enforce continuity of the displacement approximation, and local-in-space-and-time absorbing boundaries yields a second-order-accurate in time and space method that is capable of scaling up to the very large problem sizes that are required for high resolution earthquake modeling. In the next section we discuss incorporation of anelastic attenuation.

4 Anelastic Attenuation

Many studies have found that anelastic attenuation plays an important role in earthquake-induced ground motion modeling in sedimentary basins such as Los Angeles because elastic waves may be overamplified by trapped waves that are reflected and refracted within basins [Frankel and Vidale (1992); Yomogida and Etgen (1993); Olsen and Archuleta (1996); Olsen, Nigbor, and Konno (2000); Wald and Graves (1998); Pitarka, Irikura, Iwata, and Sekiguchi (1998); Sato, Graves, and Somerville (1999)]. There are several difficulties in incorporating attenuation within time domain seismic wave propagation simulations, including the large memory necessary to store additional variables that are usually required to represent anelastic attenuation, and the difficulty in estimating the quality factor (Q) for highly heterogeneous media. Day and Bradley (2001) have suggested a coarse-grained implementation of the memory variables that improves performance in terms of accuracy of solution and efficiency of memory use. Moreover, Olsen, Day, and Bradley (2003) have demonstrated an effective Q model for the LA basin in the context of simulations of the 1994 Northridge earthquake.

In this study, we instead use a Rayleigh damping model because of its convenience, simplicity, and consistency with our existing seismic wave propagation code. Here, the element damping matrix is taken proportional to element mass and stiffness matrices,

$$\mathbf{C}^e = \kappa_1 \mathbf{M}^e + \kappa_2 \mathbf{K}^e, \quad (6)$$

where $\mathbf{C}^e$ is the element material damping matrix, $\mathbf{M}^e$ is the element mass matrix, $\mathbf{K}^e$ is the element stiffness matrix, and κ_1 and κ_2 are scalar coefficients that are to be determined element-wise with the goal that the damping ratio for each mode is independent of the frequency over a given frequency range. Pre- and post-multiplying (6) by the n th eigenvector,

$$\phi_n^T \mathbf{C}^e \phi_n = \kappa_1 \phi_n^T \mathbf{M}^e \phi_n + \kappa_2 \phi_n^T \mathbf{K}^e \phi_n, \quad (7)$$

and making use of orthonormality and the modal damping assumption, we obtain

$$2\omega_n \zeta_n = \kappa_1 + \kappa_2 \omega_n^2 \quad (8)$$

or

$$\zeta_n = \frac{1}{2\omega_n} \kappa_1 + \frac{\omega_n}{2} \kappa_2, \quad (9)$$




Figure 3 : Relation between damping ratio and shear wave velocity

where the damping ratio ζ_n is equivalent to $\frac{1}{2Q}$ and ω_n is the natural frequency associated with the n th mode. Assuming that the damping ratio ζ varies with the inverse of the shear wave velocity, β [as in Bao, Bielak, Ghattas, Kallivokas, O'Hallaron, Schewchuk, and Xu (1998)],

$$\zeta = \frac{\gamma_1}{\beta + \gamma_2}, \quad (10)$$

where γ_1 and γ_2 are scalars determined by assuming that ζ is 2% at the shear velocity β of 200 m/s while ζ is 0.5% at β of 1000 m/s. This yields $\gamma_1 = 5.333$ and $\gamma_2 = 66.67$, for which the relation between the damping ratio and the shear wave velocity is shown in Figure 3. For any given shear velocity, then, (10) determines the desired damping ratio. Our goal is to find the Rayleigh constants κ_1 and κ_2 such that (9) yields as close to this desired damping ratio as possible over a realistic frequency range (ω_1, ω_2). This can be formulated as the linear least squares problem:

$$\min_{(\kappa_1, \kappa_2)} \int_{\omega_1}^{\omega_2} \left(\zeta - \frac{\kappa_1}{2\omega} - \frac{\kappa_2 \omega}{2} \right)^2 d\omega \quad (11)$$

For example, solving this least squares problem for $\zeta = 0.5\%$ over the frequency range (0.1, 1.1) Hz gives the plot in Figure 4 of ζ as determined by (9). The near-constant ζ , and thus Q , over the frequency range of interest, is evident in the figure. The Rayleigh damping model, which




Figure 4 : Relation between damping ratio and frequency for an optimized Rayleigh damping model

increases both linearly and inversely with frequency, provides a reasonable damping model for many soils, although very low and very high frequencies are overdamped.

5 Numerical Tests for Validation

To verify our octree-based seismic wave propagation code, we have conducted a number of numerical tests ranging from idealized models to realistic basins with various properties and excitations Kim (2003). Here, we present four representative examples. In all cases, the mesh is generated with the out-of-core octree/hexahedral mesh generator Euclid [Tu, O’Hallaron, and López (2002); Tu and O’Hallaron (2004)], and partitioned with the parallel mesh partitioner ParMETIS [Karypis, Schloegel, and Kumar (2002)].

First, we solve a layered half space problem with point source excitation. The problem description is given in Table 1. Figure 5a illustrates the location of the point source excitation.

The medium is piecewise-homogeneous, consisting of a homogeneous elastic layer over an elastic half space. The excitation is defined as a double-couple, located in the half space as a point source. The mesh size is chosen so that shear wavelengths corresponding to a frequency of 1 Hz are resolved with 10 nodes per wavelength. This results in an element size in the half space of 450 m

Table 1 : Layered half space problem description

domain size:	$36 \times 36 \times 18 \text{ km}^3$
depth of layer:	1.8 km
layer:	$\rho = 2500 \text{ kg/m}^3$ $\alpha = 3900 \text{ m/s}$ $\beta = 2250 \text{ m/s}$
half space:	$\rho = 3048 \text{ kg/m}^3$ $\alpha = 7800 \text{ m/s}$ $\beta = 4500 \text{ m/s}$
excitation:	$M_0[\frac{t}{T_0} - \frac{1}{2\pi}\sin(2\pi\frac{t}{T_0})]$ $0 \leq t \leq T_0$ M_0 $t > T_0$
strike, dip, rake:	$30^\circ, 40^\circ, 60^\circ$
magnitude (M_0):	$1.4 \times 10^{13} \text{ Nm}$
rise time (T_0):	1 s
source location:	$(-225 \text{ m}, 225 \text{ m}, -6075 \text{ m})$

and in the top layer of 225 m, resulting in a mesh with 435,200 elements, 469,485 total nodes, and 19,360 hanging nodes. The time step size is 0.02 s for the 40 s simulation. A Cartesian coordinate system is chosen with x positive in the East (E) direction, y positive in the North (N) direction, and z positive in the Upward (U) direction. This axis convention applies to the remaining examples. Figures 5b, 5c, and 5d provide a comparison of the octree FEM solution with that from an analytical Green’s function code.⁵ We compare at 8 observation points near the top surface along a diagonal from the epicenter. Both numerical and analytical results are low-pass filtered to 1 Hz, the target resolution. These results show good graphical agreement between the two solutions.

To investigate the influence of mesh irregularity due to octree refinement, we construct a mesh that is adapted to the shear wave velocities of the Southern California Earthquake Center (SCEC) model version 2.2 [Magistrale, Day, Clayton, and Graves (2000)] for a frequency of 0.2 Hz. We replace the spatially-variable SCEC velocity model with a homogeneous material model in order to compare with the Green’s function solution. Table 2 lists the problem data. Figure 6 depicts the irregular mesh for this model. The ratio of the largest element size ($h = 1250 \text{ m}$) and smallest element size ($h = 156.25 \text{ m}$) is 8, corresponding to an octree with four levels. Comparisons with Green’s function solutions are shown in Figure 7a, 7b, and 7c for a number of observation points that are

⁵ The Green’s function code was written by Y. Hisada.




Figure 5 : Comparison of octree (dotted) with Green's function (solid) solution for layered half space problem. (a) Problem description. (b) E-W component of velocity. (c) N-S component of velocity. (d) U-D component of velocity.

located in a similar pattern as in the half space problem. Results are filtered to 0.2 Hz. The comparison with the analytical solution is again very good, suggesting correct treatment of the mesh size transitions.

To examine the implementation of anelastic attenuation, we use the same model that was used for the layered half space problem (Figure 5a). We choose damping ratios $\zeta_1 = 15\%$ ($Q_s = Q_p = 3.33$) for the layer and $\zeta_2 = 3\%$ ($Q_s = Q_p = 16.67$) for the half space. The damping ra-

tios are high for the given properties; in fact our goal is to verify the anelastic implementation with the Green's function solution for large attenuation. This comparison is shown in Figure 8a, 8b, and 8c and again the agreement is very good.

Finally, we perform a simulation of the 1994 Northridge earthquake in the LA Basin and compare our octree-based FEM results with synthetic seismograms from the FDM code of Graves (1998) and our earlier tetrahe-

Table 2 : Half space with irregular octree mesh

domain size:	$80 \times 80 \times 30 \text{ km}^3$
half space:	$\rho = 2000 \text{ kg/m}^3$ $\alpha = 5000 \text{ m/s}$ $\beta = 2500 \text{ m/s}$
excitation:	$M_0[\frac{t}{T_0} - \frac{1}{2\pi}\sin(2\pi\frac{t}{T_0})] \quad 0 \leq t \leq T_0$ $M_0 \quad t > T_0$
strike, dip, rake:	$30^\circ, 40^\circ, 60^\circ$
magnitude (M_0):	$1.4 \times 10^{13} \text{ Nm}$
rise time (T_0):	1 s
source location:	(40.625 km, 40.625 km, -14.375 km)




**Figure 7** : Comparison of octree (dotted) with Green's function (solid) solution for irregular mesh of LA Basin, with homogeneous properties. (a) E-W component of velocity. (b) N-S component of velocity. (c) U-D component of velocity.**Figure 8** : Comparison of octree (dotted) with Green's function (solid) solution for anelastic model of layered half space problem. (a) E-W component of velocity. (b) N-S component of velocity. (c) U-D component of response.

Table 3 : Performance analysis of octree code for 1994 Northridge earthquake simulation

Model	0.1 Hz	0.2 Hz	0.5 Hz	1.0 Hz	1.0 Hz	1.0 Hz
Procs	1	16	128	512	1024	2048
Nodes ($\times 10^6$)	0.134	0.618	14.8	47.6	102	102
MFLOPS/s	505	491	469	451	450	443
Par Effcn	100%	97%	93%	89%	89%	88%
CPU Effcn	25.3%	24.6%	23.5%	22.6%	22.5%	22.2%

**Figure 6** : Irregular mesh for LA basin model

dral code [Bao, Bielak, Ghattas, Kallivokas, O'Hallaron, Shewchuk, and Xu (1996); Bao, Bielak, Ghattas, Kallivokas, O'Hallaron, Schewchuk, and Xu (1998)]. The dimensions of the basin model are $80 \times 80 \times 30 \text{ km}^3$. Material properties are taken from the SCEC velocity model and the fault rupture model is based on Wald, Heaton, and Hudnut (1996). The maximum resolved frequency is 0.5 Hz. Anelastic attenuation is based on Q factors given by Olsen, Day, and Bradley (2003). The resulting mesh has approximately 12 million elements and 15 million nodes. Figure 9 compares synthetic results from the three codes at a number of observation stations. The synthetic seismograms have been low-pass filtered using a Butterworth filter with 4 poles and 2 passes. The overall agreement once again is very good.

6 Performance and scalability

Our octree-based finite element code has achieved excellent performance and isogranular scalability over a range of 1 to 2048 processors on the HP AlphaServer Cluster

at the Pittsburgh Supercomputing Center. Table 3 provides details on the performance of the code for a series of simulations of the 1994 Northridge earthquake in the LA basin corresponding to highest resolved frequency ranging from 0.1 Hz (run on 1 processor) to 1 Hz (run on 2048 processors). Problem size varies from 134,500 to over 100 million nodes. Scalability is excellent: 88% parallel efficiency is maintained in scaling from 1 to 2048 processors. The scalar efficiency is also very good, degrading to just 22.2% of peak on each Alpha node (based on 2 GFLOPS/s peak). This latter figure is very good for a modern microprocessor, considering the irregular nature of the computations. The 2048 processor simulation sustains nearly a TFLOPS/s.

7 Concluding Remarks

We have presented a parallel octree-based finite element method for large-scale earthquake ground motion simulation in realistic basins. The octree representation combines the low memory per node and good cache performance of finite difference methods with the spatial adaptivity to local seismic wavelengths characteristic of unstructured finite element methods. Several tests are provided to verify the numerical performance of the method against Green's function solutions for homogeneous and piecewise homogeneous media, both with and without anelastic attenuation. A comparison is also provided against a finite difference code and an unstructured tetrahedral finite element code for a simulation of the 1994 Northridge Earthquake. The numerical tests all show very good agreement with analytical solutions and other codes. Finally, performance evaluation indicates excellent single-processor performance and parallel scalability over a range of 1 to 2048 processors for Northridge simulations with up to 300 million degrees of freedom.

Acknowledgement: Two words stand out in our




Figure 9 : Comparison of synthetic seismograms at 11 stations for 1994 Northridge Earthquake simulation. Top seismogram corresponds to FDM simulation; middle to tetrahedral FEM; bottom to octree FEM.

minds when we think of Cliff Astill: vision and integrity. Soon after moving to his new position as Program Director of Ground Motion and Geotechnical Earthquake Engineering at NSF in the late 1980s, Cliff came to the realization that, along with experimentation and direct field observation, model-based simulation would have to play a leading role in research, along with experimentation and direct field observation, if real advances were to be made in the understanding and prediction of the earthquake behavior of complex geotechnical systems. In his unassuming, yet highly effective style, he was instrumental in creating VELACS, a multi-investigator project for the validation of liquefaction analysis using centrifuge studies. This was the first large-scale concerted effort to evaluate the utility of numerical codes that included pore pressure development and liquefaction modeling abilities to function as real-life design codes. Through this and other activities, Cliff became one of the primary forces in setting the course of geotechnical earthquake engineering for the next generation.

In our own work, Cliff provided a willing ear, and initially his own discretionary funds, to support our exploratory work on what would later become a major effort on earthquake ground motion simulation in large basins. Even though we were neophytes at the time, with his wisdom and encouragement he made us feel from the very beginning as if we knew more than we did, thereby giving us the confidence to actually achieve something useful. Most appropriately, Cliff later served as NSF cognizant officer for both the Grand Challenges (CMS-9318163) and KDI (CMS-9980063) projects. The latter supported the bulk of the underlying work for the present paper.

While our initial relationship was purely professional, over the years it grew into a real friendship. Through it, we came to know Cliff as a humanist, lover of literature, and great athlete. Never would we have imagined that he was a whitewater canoeist of the highest caliber or that he would regularly participate with his wife in walkathons that lasted more than 20 hours.

The illness that would ultimately take his life made apparent the strength of character that we already knew well. Always having had a keen intellectual curiosity, Cliff became an expert on his own disease, and was able to participate in his own treatment, and sometimes guide his physicians in the appropriate course of action. Even at the stage when many might have despaired, Cliff lived

life at its fullest to the very end.

For us, it is a privilege to have had the opportunity of knowing and working with Cliff Astill, a first-rate scholar, kind and generous person, gentleman, and a wonderful human being. We dedicate this paper to his memory.

We would also like to express our appreciation to George K. Lea, who himself has provided guidance and encouragement to us through the years, for his kind invitation to submit a contribution to this special issue of the journal.

Partial support for this work was also provided through grants from the NSF-ITR program (EAR-0326449) and the Southern California Earthquake Center. SCEC is funded by NSF Cooperative Agreement EAR-0106924 and USGS Cooperative Agreement 02HQAG0008. The SCEC contribution number for this paper is 947. We also thank the Pittsburgh Supercomputing Center for its continued support of our large scale earthquake modeling research. In particular, the simulations in this work were conducted on the HP AlphaServer parallel system at PSC under awards BCS020001P and MCA01S002P.

References

- Akcelik, V.; Bielak, J.; Biros, G.; Epanomeritakis, I.; Fernandez, A.; Ghattas, O.; Kim, E. J.; Lopez, J.; O'Hallaron, D. R.; Tu, T.; Urbanic, J. (2003): High-resolution forward and inverse earthquake modeling on terascale computers. In *Proceedings of ACM/IEEE SC2003*, Phoenix, AZ.
- Aki, K.; Richards, P. G. (1980): *Quantitative Seismology: Theory and Methods*, volume I. W. H. Freeman and Co.
- Bao, H.; Bielak, J.; Ghattas, O.; Kallivokas, L.; O'Hallaron, D. R.; Schewchuk, J. R.; Xu, J. (1998): Large-scale simulation of elastic wave propagation in heterogeneous media on parallel computers. *Computer Methods in Applied Mechanics and Engineering*, vol. 152, pp. 85–102.
- Bao, H.; Bielak, J.; Ghattas, O.; Kallivokas, L. F.; O'Hallaron, D. R.; Shewchuk, J. R.; Xu, J. (1996): Earthquake ground motion modeling on parallel computers. In *Supercomputing '96*, Pittsburgh, Pennsylvania.

- Day, S. M.; Bradley, C. R.** (2001): Memory-efficient simulation of anelastic wave propagation. *Bulletin of the Seismological Society of America*, vol. 91, pp. 520–531.
- Dravinski, M.** (1982): Influence of interface depth upon strong ground motion. *Bulletin of the Seismological Society of America*, vol. 72, pp. 596–614.
- Frankel, A.; Vidale, J.** (1992): A three-dimensional simulation of seismic waves in the Santa Clara Valley, California from a Loma Prieta aftershock. *Bulletin of the Seismological Society of America*, vol. 82, pp. 2045–2074.
- Graves, R. W.** (1996): Simulating seismic wave propagation in 3D elastic media using staggered-grid finite differences. *Bulletin of the Seismological Society of America*, vol. 86, pp. 1091–1106.
- Graves, R. W.** (1998): Three-dimensional finite-difference modeling of the San Andreas fault: Source parameterization and ground motion levels. *Bulletin of the Seismological Society of America*, vol. 88, pp. 891–897.
- Karypis, G.; Schloegel, K.; Kumar, V.** (2002): ParMETIS 3.0. <http://www-users.cs.umn.edu/~karypis/metis/parmetis/>, 2002.
- Kawase, H.** (1988): Time domain response of a semi-circular canyon for incident SV, P, and Rayleigh waves calculated by the discrete wavenumber boundary element method. *Bulletin of the Seismological Society of America*, vol. 78, pp. 1415–1437.
- Kim, E. J.** (2003): *Parallel octree-based multiresolution mesh method for large scale earthquake ground motion simulation*. PhD thesis, Department of Civil and Environmental Engineering, Carnegie Mellon University, 2003.
- Kim, E. J.; Bielak, J.; Ghattas, O.** (2003): Large-scale Northridge Earthquake simulation using octree-based multiresolution mesh method. In *16th ASCE Engineering Mechanics Conference*.
- Komatitsch, D.; Tromp, J.** (1999): Introduction to the spectral-element method for 3-D seismic wave propagation. *Geophysics Journal International*, vol. 139, pp. 806–822.
- Levander, A. R.** (1988): Fourth-order finite-difference P-SV seismograms. *Geophysics*, vol. 53, pp. 1425–1436.
- Li, X.; Bielak, J.; Ghattas, O.** (1992): Three-dimensional earthquake site response on a CM-2. In *Proceedings of the 10th World Conference on Earthquake Engineering*.
- Magistrale, H.; Day, S.; Clayton, R. W.; Graves, R.** (2000): The SCEC Southern California reference three-dimensional seismic velocity model version 2. *Bulletin of the Seismological Society of America*, vol. 90, pp. S65–S76.
- Olsen, K. B.; Archuleta, R. J.** (1996): Three-Dimensional simulation of earthquakes on the Los Angeles fault system. *Bulletin of the Seismological Society of America*, vol. 86, pp. 575–586.
- Olsen, K. B.; Day, S. M.; Bradley, C. R.** (2003): Estimation of Q for long-period (> 2 s) waves in the Los Angeles Basin. *Bulletin of the Seismological Society of America*. In press.
- Olsen, K. B.; Nigbor, R.; Konno, T.** (2000): 3-D viscoelastic wave propagation in the upper Borrego Valley, California constrained by borehole and surface data. *Bulletin of the Seismological Society of America*, vol. 90, pp. 134–150.
- Pitarka, A.; Irikura, K.; Iwata, T.; Sekiguchi, H.** (1998): Three-dimensional simulation of the near fault ground motion for the 1995 Kyogo-ken Nanbu (Kobe), Japan, earthquake. *Bulletin of the Seismological Society of America*, vol. 88, pp. 428–440.
- Sánchez-Sesma, F. J.** (1978): Ground motion amplification due to canyons of arbitrary shape. In *Proceedings of 2nd International Conference on Microzonation*, pp. 729–738, San Francisco, California.
- Sánchez-Sesma, F. J.; Bravo, M. A.; Herrera, I.** (1985): Surface motion of topographical irregularities for incident P, SV, and Rayleigh waves. *Bulletin of the Seismological Society of America*, vol. 75, pp. 263–269.
- Sánchez-Sesma, F. J.; Esquivel, J. A.** (1979): Ground motion on alluvial valley under the incident plane SH waves. *Bulletin of the Seismological Society of America*, vol. 69, pp. 1107–1120.

- Sato, T.; Graves, R. W.; Somerville, P. G.** (1999): Three-dimensional finite-difference simulations of long-period strong motions in the Tokyo metropolitan area during the 1990 Odawara Earthquake (Mj 5.1) and the Great 1923 Kanto earthquake (Ms 8.2) in Japan. *Bulletin of the Seismological Society of America*, vol. 89, pp. 579–607.
- Schriener, C. W.; Helmberger, D. V.** (1994): Seismic waveform modeling in the Los Angeles Basin. *Bulletin of the Seismological Society of America*, vol. 84, pp. 1310–1326.
- Toshinawa, T.; Ohmachi, T.** (1988): Love-wave propagation in a three-dimensional basin. *Bulletin of the Seismological Society of America*, vol. 82, pp. 1661–1677.
- Tu, T.; O'Hallaron, D.; López, J.** (2002): Etree: A database-oriented method for generating large octree meshes. In *Proceedings of the Eleventh International Meshing Roundtable*.
- Tu, T.; O'Hallaron, D. R.** (2004): A computational database system for generating unstructured hexahedral meshes with billions of elements. In *Proceedings of ACM/IEEE SC2004*, Pittsburgh, PA.
- Vidale, J. E.; Helmberger, D. V.** (1988): Elastic finite-difference modeling of the 1971 San Fernando, California, earthquake. *Bulletin of the Seismological Society of America*, vol. 78, pp. 122–141.
- Virieux, J.** (1984): SH-wave propagation in heterogeneous media: Velocity-stress finite-difference method. *Geophysics*, vol. 49, pp. 1933 – 1957.
- Wald, D. J.; Graves, R. W.** (1998): The seismic response of the Los Angeles Basin, California. *Bulletin of the Seismological Society of America*, vol. 88, pp. 337–356.
- Wald, D. J.; Heaton, H.; Hudnut, K. W.** (1996): The slip history of the 1994 Northridge, California, Earthquake determined from strong-motion, teleseismic, GPS, and leveling data. *Bulletin of the Seismological Society of America*, vol. 86, pp. S49–S70.
- Wong, H. L.** (1979): Diffraction of P, SV, and Rayleigh waves by surface topographies. Technical Report Report CE79-05, Department of Civil Engineering, University of Southern California, 1979.
- Wong, H. L.** (1982): Effect of surface topography on the diffraction of P, SV, and Rayleigh waves. *Bulletin of the Seismological Society of America*, vol. 72, pp. 1167–1183.
- Yomogida, K.; Etgen, J. T.** (1993): 3-D wave propagation in the Los Angeles basin for the Whittier-Narrows earthquake. *Bulletin of the Seismological Society of America*, vol. 83, pp. 1325–1344.
- Yoshimura, C.; Bielak, J.; Hisada, Y.; Fernandez, A.** (2003): Domain reduction method for three-dimensional earthquake modeling in localized regions, Part II: Verification and application. *Bulletin of the Seismological Society of America*, vol. 93, pp. 825 – 840.
- Young, D.; Melvin, R.; Bieterman, M.; Johnson, F.; Samant, S.; Bussoletti, J.** (1991): A locally refined rectangular grid finite element method: application to computational fluid dynamics and computational physics. *Journal of Computational Physics*, vol. 92, pp. 1–66.
- Zang, L.; Chopra, A.** (1991): Three dimensional analysis of spatially varying ground motion around uniform canyon in a homogeneous half space. *Earthquake Engineering and Structural Dynamics*, vol. 20, pp. 911–926.