5B.4

CUBIT and Seismic Wave Propagation Based Upon the Spectral-element Method: An Advanced Unstructured Mesher for Complex 3D Geological Media

Emanuele Casarotti[1,5], Marco Stupazzini[2], Shiann Jong Lee[3], Dimitri Komatitsch[4], Antonio Piersanti[5], and Jeroen Tromp[6]

[1] Seismological Laboratory, California Institute of Technology. 1200 E. California Blvd., MS 252-21, Pasadena, California, 91125, USA. emanuele@gps.caltech.edu

[2] Department of Earth- and Environmental Sciences, Geophysics section, Ludwig-Maximilians-Universität. Theresienstrasse 41, 80333, Munich, Germany. stupa@geophysik.uni-muenchen.de

[3] Institute of Earth Science, Academia Sinica. No.128, Section2, Academia Road, Nankang, Taipei, Taiwan 115. sjlee@earth.sinica.edu.tw

[4] Laboratoire de Modlisation et d'Imagerie en Gosciences UMR 5212, Universit de Pau et des Pays de l'Adour. Btiment IPRA - Avenue de l'Universit, BP 1155, 64013, Pau, France. dimitri.komatitsch@univ-pau.fr

[5] Istituto Nazionale di Geofisica e Vulcanologia. Via di Vigna Murata 605, 00143, Rome, Italy. piersanti@ingv.it

[6] Seismological Laboratory, California Institute of Technology. 1200 E. California Blvd., MS 252-21, Pasadena, California, 91125, USA. jtromp@gps.caltech.edu

Abstract

Unstructured hexahedral mesh generation is a critical part of the modeling process in the Spectral-Element Method (SEM). We present some examples of seismic wave propagation in complex geological models, automatically meshed on a parallel machine based upon CUBIT (Sandia Laboratory, cubit.sandia.gov), an advanced 3D unstructured hexahedral mesh generator that offers new opportunities for seismologist to design, assess, and improve the quality of a mesh in terms of both geometrical and numerical accuracy. The main goal is to provide useful tools for understanding seismic phenomena due to surface topography and subsurface structures such as low wave-speed sedimentary basins. Our examples cover several typical geophysical problems: 1) "layer-cake" volumes with high-resolution topography and complex solid-solid interfaces (such as the Campi Flegrei Caldera Area in Italy), and 2) models with an embedded sedimentary basin (such as the Taipei basin in Taiwan or the Grenoble Valley in France).

1 The spectral-element method in seismology

Wave propagation phenomena can nowadays be studied thanks to many powerful numerical techniques. We have seen rapid advances in computational seismology at global, regional, and local scales thanks to various numerical schemes, such as the finite-difference method (FDM) [15], the finite-element method (FEM) [1], the spectral-element method (SEM) [12], the integral boundary-element method (IBEM) [2], and the Arbitrary High-Order Discontinuous Galerkin method (ADER-DG) [8]. Spurred by the computational power made available by parallel computers, geoscientists and engineers can now accurately compute synthetic seismograms in realistic 3D Earth models. Among the methods previously listed, the SEM has convincingly demonstrated the ability to handle high-resolution simulations of seismic wave propagation in 3D domains.

The purpose of this article is not to give a complete description of the SEM, but a basic introduction to the main properties of the SEM is needed in order to understand the various constraints imposed on the meshing process. The SEM is as a generalization of the FEM based on the use of high-order piecewise polynomial functions. A crucial aspect of the method is its capability to provide an arbitrary increase in accuracy by simply increasing the algebraic degree of these functions (the spectral degree n). From a practical perspective, this operation is completely transparent to the users, who limit themselves to choosing the spectral degree at runtime, leaving to the software the task of building up suitable quadrature points and the relevant degrees of freedom. Obviously, one can also use mesh refinement to improve the accuracy of the numerical solution, thus following the traditional finite-element approach.

Referring to the literature for further details [12], we begin by briefly summarizing the key features of the SEM. We start from the general differential form of the equation of elastodynamics:

$$\rho \partial_t^2 \mathbf{s} = \nabla \cdot \mathbf{T} + \mathbf{F}, \tag{1}$$

where $\mathbf{s}(\mathbf{x}, t)$ denotes the displacement at position $\mathbf{x}$ and time t, $\rho(\mathbf{x})$ is the density distribution, and $\mathbf{F}(\mathbf{x}, t)$ the external body force. The stress tensor $\mathbf{T}$ is related to the strain tensor by the constitutive relation $\mathbf{T} = \mathbf{c} : \nabla \mathbf{s}$, where $\mathbf{c}$ denotes a fourth-order tensor. No particular assumptions are made regarding the structure of $\mathbf{c}$, which describes the (an)elastic properties of the medium (the formulation is general and can incorporate full anisotropy [9] or non-linear constitutive relationships [6]).

For seismological applications, a numerical technique needs to facilitate at least the following: (i) an attenuating medium, (ii) absorption of seismic energy on the fictitious boundaries of the domain in order to mimic a semi-infinite medium (the free-surface condition is a natural condition in the SEM), and finally (iii) seismic sources. In the SEM all these features are readily accommodated [12].

As in the FEM, the dynamic equilibrium problem for the medium can be written in a weak, or variational, form, and, through a suitable discretization procedure that depends on the numerical approach, can be cast as a system of ordinary differential equations with respect to time. Time-marching of this system of equations may be accomplished based upon an explicit second-order finite-difference scheme, which is conditionally stable and must satisfy the well known Courant condition [4].

The key features of the SEM discretization are as follows:

1. Like in the FEM, the model volume Ω is subdivided into a number of non-overlapping elements $\Omega_e, e = 1, \ldots, n_e$, such that $\Omega = \bigcup_{e=1}^{n_e} \Omega_e$.
2. The expansion of any function within the elements is accomplished based upon Lagrange polynomials of suitable degree n constructed from $n + 1$ interpolation nodes.
3. In each element, the interpolation nodes are chosen to be the Gauss-Lobatto-Legendre (GLL) points, i.e., the $n+1$ roots of the first derivatives of the Legendre polynomial of degree n. On such nodes, the displacement, its spatial derivatives, and integrals encountered in the weak formulation are evaluated.
4. The spatial integration is performed based upon GLL quadrature (while most classical FEMs use Gauss quadrature).

Thanks to this numerical strategy, exponential accuracy of the method is ensured and the computational effort minimized, because the resulting mass matrix is exactly diagonal. The latter feature does not occur in so-called hp FEMs, nor in SEMs based upon Chebychev polynomials.

2 Mesh design for Spectral-Element Methods

The first critical ingredient required for a spectral-element simulation is a high-quality mesh appropriate for the 3D model of interest. This process generally requires discouraging expertise in meshing and preprocessing, and is subject to several tough constraints: 1) the number of grid points per shortest desired wavelength, 2) the numerical stability condition, 3) an acceptable distortion of the elements, 4) balancing of numerical cost and available computing resources. A poor-quality mesh can generate numerical problems that lead to an increase in the computational cost, poor (or lack of) convergence of the simulation, or inaccurate results. For example, a geological model often includes a layered volume, and a staircase sampling of the interfaces between the layers can produce fictitious diffractions. Therefore, a good mesh should honor at least the major geological discontinuities of the model.

As noted in Section 1, the SEM is similar to a high-degree FEM, and in fact these methods share the first part of the discretization process. The present paper is focused on this first meshing step, and thus the results are relevant to both FEMs and SEMs. In the SEM, the subsequent step is the

evaluation of the model on the GLL integration points [12]. Here we only note that if we use Lagrange polynomials and GLL quadrature, the mass matrix is exactly diagonal, resulting in a dramatic improvement in numerical accuracy and efficiency. With this choice, each element of the mesh contains $(n+1)^3$ GLL points. Unfortunately, this approach requires that the mesh elements are hexahedra [12]. It is worth mentioning that for 2D problems it is possible to develop a SEM with triangles keeping the diagonal mass matrix but with an higher numerical cost [10]. The fact that we are restricted to hexahedral elements complicates matters significantly. Whereas 3D unstructured tetrahedral meshes can be constructed relatively easily with commercial or non-commercial codes, the creation of 3D unstructured hexahedral meshes remains a challenging and unresolved problem. For complex models, as in the case of realistic geological volumes, generating an all-hexahedral mesh based upon the available meshing algorithms can require weeks or months, even for an expert user [21]. One of the features of the SEM that impacts the creation of the mesh is the polynomial degree n used to discretize the wave field. The following heuristic rule has emerged to select n for an unstructured mesh of a heterogeneous medium: if $n < 4$ the inaccuracies are similar to the standard FEM, while if $n > 10$ the accuracy improves but the numerical cost of the simulation becomes prohibitive. The choice of n is related to the grid spacing Δh: in order to resolve the wave field to a shortest period T_0, the number of points per wavelength λ should be equal or greater than 5, leading to the constraint expressed in Eq. 2. If $n = 4$, then Δh is roughly equal to λ. We note that, for the same accuracy, a classical low-degree FEM requires a higher number of elements. Since the material properties are stored for each GLL point and can vary inside an element, we are able to interpolate the geological interfaces that our mesh is not able to follow. Nevertheless, this is an undesirable staircase sampling of the model, which introduces non-physical diffractions. Therefore, it is necessary that the mesh honors the major seismological contrasts. Furthermore, when a discontinuity is honored, the elements across the interface share some nodes, which will have the properties of the material below the discontinuity in one case, and the proprieties of the material above the discontinuity in the other case. Thus, the actual behavior of the seismic waves at the geological interface is perfectly mimicked in a way that cannot be achieved by an interpolation solely based upon Lagrange polynomials and the GLL quadrature.

Another constraint on the design of the mesh is the stability condition imposed by the adoption of an explicit conditionally stable time scheme. For a given mesh, there is an upper limit on the time step above which the calculations are unstable. We define the Courant stability number of the time scheme as $C = \Delta t(v/\Delta h)_{\max}$, where Δt is the time step and $(v/\Delta h)_{\max}$ denotes the maximum ratio of the compressional-wave speed and the grid spacing. The Courant stability condition ensures that calculations are stable if the Courant number is not higher than an upper limit $C_{\max}$ [4], providing a condition that

determines the time step Δt (Eq. 3). Again, an heuristic rule suggests that $C_{\max}$ is roughly equal 0.3–0.4 for a deformed and heterogeneous mesh [12].

Like any technique based upon an unstructured deformed mesh, the SEM requires a smooth variation of the Jacobian and an acceptable distortion of the elements. Usually, to produce an acceptable accuracy the maximum equiangle skewness should not be greater than 0.5, although empirical tests show that $s < 0.8$ can sometimes be acceptable (Eq. 4) [12].

To sum up, spectral-element simulations require an unstructured all-hexahedral conforming mesh subject to the following heuristic constraints:

$$\Delta h = v_{\min} T_0 \frac{n+1}{f(n)}, \tag{2}$$

$$\Delta t < C_{\max} \frac{v_{\min}}{v_{\max}} T_0 \frac{n+1}{f(n)}, \tag{3}$$

$$s < 0.8, \tag{4}$$

where $v_{\max}$ and $v_{\min}$ are the maximum and minimum seismic wave speeds in the element, $\mathrm{C}_{\max}$ is the Courant constant discussed above, T_0 is the shortest period that we seek to resolve, and $f(n)$ is an empirical function related to the number of points per wavelength and the polynomial degree n; for $n = 4$ $f(n) = 5$.

The mesh should have a similar number of grid points per wavelength throughout the entire model (i.e., the same numerical resolution everywhere), and since seismic wave speeds vary inside the volume, the mesh should have a variable element density, in accordance with Eq. 1. Wave speeds increase with depth in a typical geological model, so that the mesh should be denser near the surface and close to low wave-speed regions, such as an alluvial basin. In elements located at the top surface, $v_{\min}$ is the surface-wave speed, which controls the resolution of the mesh. Inside the volume it is the shear waves that have the slowest speed and thus determine Δh. From Eq. 3, we note that the smaller the grid spacing, the smaller the time step needed to obtain numerical stability, at the expense of computational cost. Therefore, Δh should be carefully evaluated in order to optimize numerical efficiency and accurate geometrical representation of the model. In the current SEM implementation, the time step is constant for the whole model. It is evident that a mesh that coarsens with depth accommodates a larger time step than a mesh with a constant element size throughout the entire volume, leading to a significant saving in computational cost.

Finally, one additional difficulty is the enormous number of elements that a typical seismic wave propagation simulation requires, ranging from hundreds of thousands of elements for local-scale simulations with low-frequency content, to tens of millions of elements for regional-scale simulations or for a high-frequency local-scale simulation. This implies that we need a parallel computational approach not only for the simulation of seismic wave propagation, but also for the mesh generation process.

The package of choice for our simulations is SPECFEM3D (`geodynamics.org`), which simulates global, regional (continental-scale) and local seismic wave propagation. Effects due to lateral variations in compressional-wave speed, shear-wave speed, density, a 3-D crustal model, ellipticity, topography and bathymetry, the oceans, rotation, and self-gravitation are all included, as well as full 21-parameter anisotropy and lateral variations in attenuation. We also present some SEM results obtained based upon the GeoELSE software package [21]. This SEM implementation is more suited for local- or small-scale simulations of seismic wave propagation (typically the goal is the study of soil-structure interaction problems). The main features of this kernel are the capability of 1) dealing with externally created 3D unstructured meshes, 2) handling the partitioning and load balancing of the computational domain by incorporating the METIS software library [13, 21], 3) implementing complex constitutive behavior [6] and 4) communicating with other codes through a sub-structuring interface based upon a domain reduction method [7].

Mesh examples

In this section we present some examples of unstructured hexahedral meshes that have been developed for simulations with SPECFEM3D.

The SPECFEM3D_GLOBE software package has been created primarily to simulate 3D seismic wave propagation on a global scale. For this purpose, a unique optimized mesh has been generated based upon a cubed sphere decomposition [17]. The details are described in [12]. Here we note that the need to densify the mesh close to the surface has been resolved by introducing doubling layers. The mesh honors topography and bathymetry and the major internal discontinuities, but not the Moho (the bottom of the Earth's crust) and the intra-crustal interfaces, which are interpolated based upon the Lagrange interpolation. This approximation is justified since the shapes of the Moho and the intra-crustal interfaces are insufficiently known to support the increase in elements that a full description of these boundaries would require. The mesh adequately resolves broadband data with periods down to 3.5 seconds on modern computers [12]. As an example, for a simulation executed on the Japanese Earth Simulator (JAMSTEC), the mesh was composed of 206 million hexahedral elements, corresponding to 36.6 billion degrees of freedom.

More critical from the point of view of mesh creation is the application of the SEM to regional or local problems, especially if we are interested in wave propagation with periods below 1 second. The software package SPECFEM3D_BASIN simulates seismic wave propagation in sedimentary basins, attempting to take into account the full complexity of these models. The solver is general and can be used to simulate seismic wave propagation on regional or local scales, but the mesh generator is currently specifically written for Southern California [11]. The mesh that honors the Moho, surface topography, and the deeper parts of the basement of the Los Angeles and Ventura sedimentary basins. It is coarsened twice with depth: first below the low

wave-speed layer, and then below the basement. The mechanical proprieties of the volume are stored at each GLL node, taking into account the 3D Southern California seismic model. The mesh contains approximately 700 thousand spectral elements (corresponding to 136 million degrees of freedom). The minimum resolved period of the seismic waves is roughly 2 seconds [11]. The mesh is used routinely to simulate earthquakes with a magnitude greater than 3.5 in Southern California. At periods of about 4 seconds and longer, the resulting simulated wave motion closely matches the data recorded at the stations of the Southern California Seismic Network. Animations of seismic wave propagation are posted in near real-time at `shakemovie.caltech.edu`. If we want to analyze seismic wave propagation at shorter periods, we need to refine the mesh and constrain the mesh to honor the shape of the shallow parts of the basins, which is currently not done. Furthermore, the mesher is designed for this specific region. Unfortunately, generalizing the adopted meshing procedure to other geographical areas requires experience, internal modification of the code, and a significant amount of time.

For example, the SPECFEM3D_BASIN mesh generator was modified by [14] to produce a mesh for the Taipei basin (Taiwan). This work represents the state of the art for the SEM meshing process. The SPECFEM3D_BASIN mesher was modified to take into account the notable topography around Taipei city and the complex geometry of the underlying basin. Three new control surfaces were introduced at various depths and one new doubling layer was added in order to honor the main geological features of the region. The top of the model is controlled by surface topography based upon a detailed Digital Elevation Model. Just below the surface, a buffer layer is used to accommodate mesh distortions induced by the steep topography. This is critical to ensure the accurate and stable accommodation of the free surface. This layer also defines the top of the first mesh doubling, needed to take into account the low seismic wave speeds inside the basin. The shallow sedimentary basin is introduced between the surface topography and the buffer layer. Seismic wave speeds in the basin are quite low, with a sharp contrast between the basin and the surrounding basement. Compared to the Los Angeles basin, the Taipei basin is relatively small and shallow. The meshing strategy of [11] would result in a basement layer honoring only a small portion of the deepest part of the basin, and most of the basin boundaries would be not constrained by the mesh. To overcome these limitations, the nodes were empirically adjusted close to the sedimentary basin to redefine the basin boundary mesh [14]. Several additional empirical criteria were used to define how and where to move the nodes, leading to a stable and accurate mesh implementation. The resulting mesh is composed of 4.5 million spectral elements and 900 million degrees of freedom, and it is adequate for periods down to 1 second. As in the Los Angeles basin example, we note that this meshing strategy is ad-hoc for this particular application. Furthermore, the size of the hexahedra on the surface is controlled by the speed of the surface waves inside the basin, leading to an oversampled mesh in the surrounding regions (i.e., an enormous number

of unnecessary elements, with a serious impact on the computational cost of the simulation). Finally, empirical adjustment of the nodes close to the basin basement was an attainable solution for the Taipei basin, but this approach cannot be generalized.

Meshing in parallel

The meshes discussed in the previous section are simply too large to be built on a single processor or even on a large shared memory machine: they require clusters or other parallel computers with distributed memory. The approach in SPECFEM3D is to partition the model into one mesh slice for each available CPU. Mesh generation and the corresponding wave propagation simulation are accomplished in parallel, the latter requiring only one communication phase between shared slice interfaces at each simulation time step. Since the communication map is constant throughout the simulation, it is built once and for all at the end of the mesh generation phase. The challenge is to load-balance the mesh, i.e., to make sure that each mesh slice contains roughly the same number of spectral-elements as any other mesh slice. This further limits the flexibility of the mesher, and adds further complexity to the procedure.

3 New meshing strategies: CUBIT

The aim of this article is to show examples of new approaches to the meshing process by means of the CUBIT mesh generation tool kit from Sandia National Laboratories (`CUBIT.sandia.gov`), with as an ultimate goal a more general parallel mesh generator for SPECFEM3D_BASIN. From the overview in the previous section it is clear that we are looking for software that has the following attributes: 1) the ability to accommodate general geological models, 2) minimal user intervention, 3) high resolution, 4) conforming, unstructured hexahedral elements, 5) a parallel, load-balanced implementation. The experience gained by the authors for various meshing problems [21] suggests that CUBIT comes very close to fulfilling these requirements. Although it currently lacks parallel capabilities, it incorporates a set of powerful and advanced meshing schemes developed to automatically handle the unstructured meshing problem. However, it is worth noting that meshing a large complex domain, such as any geological model, does not seem feasible with a desirable "single-button" procedure, due to the shape of the interfaces, the wide range of element sizes, and the broad range of possible geological models. A starting point would be to mesh in a parallel and automated way some basic classes of problems that are important in seismology. For example, a layer-cake volume combined with a shallow sedimentary basin.

3.1 Layer-cake geological volumes

The first step in interfacing CUBIT with the SPECFEM3D workflow is to reproduce the mesh that the current mesher creates. Since SPECFEM3D is

written in Fortran90, we have developed a Fortran code that takes advantage of the journal scripting capabilities of CUBIT. The adopted strategy is valid for all the geological models that can be described by a layer-cake volume, consisting of a stack of non-degenerate heterogeneous quasi-parallel horizontal layers. The volume is split into rectangular slices of equal size, one slice for each processor following the original MPI map of SPECFEM3D. We use a serial Fortran routine to build a journal file for each slice, creating and storing all the commands that drive the meshing process inside CUBIT. Schematically, we can summarize each journal file in five steps: 1) creation of the geological geometry by means of a bottom-up approach (i.e., vertices → splines → surfaces → volumes) and geometry healing, 2) definition of the meshing intervals along the curves based upon Eqs. 2 and 3, 3) definition of the meshing schemes and meshing, 4) refinement and smoothness, and 5) blocks and sidesets definition and export (Figure 1). One of the main advantages of using CUBIT is that the entire process can be fairly well automated both for mesh creation and quality checking. The user is only asked to provide the shape of the topographic and geological surfaces and a good choice of intervals based upon Eqs. 2 and 3. Furthermore, it is possible to add or subtract geological discontinuities and refinement layers, and to correct mesh details in an easier way than in the current workflow. Unfortunately, some constraints on the meshing scheme are needed. As we will see later, due to the conformity requirements for the whole mesh and at the MPI interfaces (i.e., the surfaces that are shared by two adjacent slices), the vertical surfaces must be structured, while the horizontal surfaces can be meshed with structured or unstructured schemes. The so-constructed volume has a 2.5D logical symmetry, therefore, the sweep algorithm [19] can be directly applied without any further decomposition. As explained in the previous section, generally the profiles of seismic wave speeds demand a mesh coarsening with depth. Since the mesh resulting from the approach described above is structured along the vertical direction, the only way to produce a vertical densification of elements is the refinement algorithm included in CUBIT [5], which produces a transition between 1 to 27 hexahedra, tripling the elements along a linear direction in only one hex sheet 1b. It is more efficient than the refinement used in [11] (i.e., a transition from 1 element to 8 in 2 hex sheets) and it is perfect for a sharp high-contrast in seismic wave speeds. Nevertheless, where the variation of seismic wave speeds is smoother, the higher transition ratio in CUBIT's refinement produces a larger difference between the number of points per wavelength across the mesh (i.e., an unevenly sampled mesh). For this reason, the possibility of introducing a refinement with a lower transition ratio is under development. In this respect, let us mention that a refinement with linear ratio of 2, gping from two elements to four in one sheet has successfully been implemented in SPECFEM3D_GLOBE V4.0. After the creation of the journal files, a second Fortran routine provides the MPI environment needed to build the mesh in parallel. CUBIT is executed on each processor (slice) and it plays back the related script file, meshing its portion of vol-

ume. A quality analysis of the skewness is then performed. Since each slice is meshed independently, we carefully check that each corresponding node in the surfaces shared by two adjacent independent slices has the same coordinates. The "refined-structured" scheme imposed in the vertical direction guarantees that the number of elements in each surface is the same; nevertheless, the possibility of a discrepancy in location between two corresponding nodes is not negligible, due to the applied smoothness algorithm and the influence of the background elements. The routine communicates the nodes to the adjacent MPI slice following the map in Figure 1c. If the nodes in the slices show a numerically significant location discrepancy compared to the corresponding neighboring nodes, they are adjusted accordingly. Then, the GLL points are evaluated for each element and the material proprieties are assigned. The GLL node numbering follows the same procedure as in [11]. Consequently, the resulting GLL mesh can be written in order to follow the format required by the SPECFEM_BASIN solver, so that simulation can be performed directly.

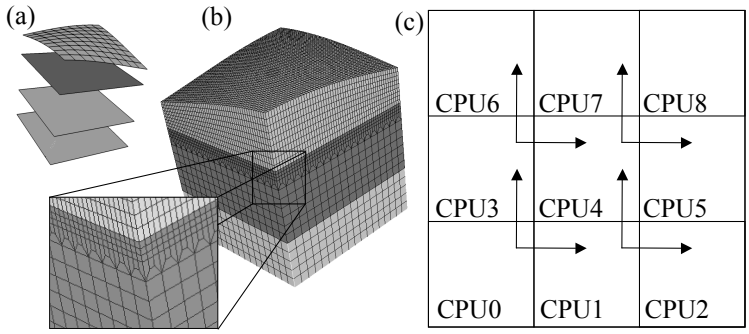

Fig. 1. Basic sketch for the mesh creation of a layer-cake geological model: (a) Step 1 of the strategy: creation of the geological geometry by means of a bottom-up approach. (b) Final result with a detailed zoom showing the tripling refinement layer in the upper part of the model. (c) The routine communicates the nodes to the adjacent MPI slices (up and right) following the map

Benchmark

In order to investigate the impact on the simulations of the coupling between SPECFEM3D and CUBIT in terms of both numerical dispersion and non-physical diffractions, we have performed some benchmarks for seismic wave propagation associated with a moderate earthquake using meshes produced by CUBIT and by the current SPECFEM3D_BASIN mesh generator. Table 1 summarizes the various benchmarks, which have been performed on 64 processors of CITERRA (a high-performance computing cluster in the Division of Geological & Planetary Sciences at Caltech [3]). The meshed volume covers part of the Taipei basin mesh described in Section 2.

Table 1. Summary of performed benchmarks

ID	model	mesh	differences
1	homogenous	regular mesh	no
2	homogenous	1 refinement	no
3	2 heterogeneous layers	1 refinement	no
4	homogenous	regular mesh & topography	negligible
5	2 layer & basin	2 refinement & topography	see Figure 2

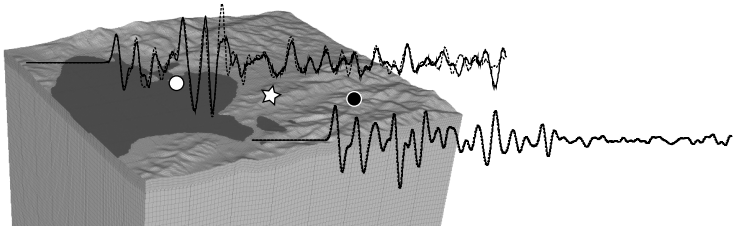

Fig. 2. Taipei basin CUBIT mesh used in the benchmark Test 5. Comparison between two seismograms obtained through two simulations of the seismic wave propagation produced by a M 6 earthquake (star): one simulation is performed with the mesh obtained by the current SPECFEM3D_BASIN mesher (dotted line) and one with the mesh created by CUBIT (solid line). The seismogram are recorded on bedrock (black circle) and soft soil basin deposits.

The goal of these benchmarks is not only to test the MPI communication capabilities of SPECFEM3D_CUBIT, but also to provide information about the impact of the tripling refinement layers and the topography on the quality of the simulation. Tests 1–3 shows that there is no significant impact due to the distorted elements in the refinement layer ($s_{\max} = 0.63$). Test 4 indicates that negligible differences in the seismograms are due to the different spline representation of the topography (Figure 1a). In Test 5 we applied the procedure for the Los Angeles basin [11] to the Taipei basin. The mesh honors the Moho interface, the topography and the deeper parts of the basin. The outcrop and the outline of the basin are interpolated by the Lagrange polynomials. Due to the different transition ratios for the refinement in SPECFEM3D_BASIN and SPECFEM3D_CUBIT, the resolution on the surface is better for SPECFEM3D_CUBIT (70 m vs 80 m). As shown in Figure 1b, if we consider the synthetic seismograms for a station on the bedrock, the two seismograms match almost perfectly, despite the different sampling. However, if we consider a station in the basin, some differences appear due to the interpolation of the discontinuity. Since no analytical solution exists for such a 3D seismological model, it is impossible to determine the seismogram that is closer to the exact solution of the problem, therefore we can only say that the results are globally in agreement for seismic waves with wavelengths greater that the dimension of the basin outcrop. Nevertheless, in order

to resolve shorter wavelengths we would need a mesh that would honor the boundary of the basin better.

Campi Flegrei Caldera

The first application of SPECFEM3D_CUBIT is to simulate seismic wave propagation in the Campi Flegrei caldera (Italy). It is an active 13-km wide quaternary volcanic caldera whose history is dominated by two eruptions that produced widespread ash-flow deposits about 34,000 years and 12,000 years ago. Intense activity has been recorded, although it is limited to small eruptions and dramatic uplift and subsidence of the terrain. Particularly noteworthy is the subsidence episode of 1982–1984 and the contemporaneous seismic crisis. The caldera is located just a few kilometers west of the city of Naples. Because of the high population density, it is considered a highly dangerous volcanic area and a challenging civil protection problem faced by authorities and scientists in case of imminent eruption.

Volcanic structures represent a challenge for seismic wave simulations due to the high level of heterogeneity and topographic variations, and the presence of fractures, magmatic intrusions, and low wave-speed regions. For our simulation, we adopted a 3D model for Campi Flegrei [18], composed of 3 layers with a P-wave speed ranging from 2.7 to 6.5 km/s. The interfaces have a resolution of 250 m, while the topography is a based upon a 60 m Digital Elevation Model. The dimensions of the volume are 40 km$\times$40 km$\times$9 km. Both the mesh generation and the simulation are performed on 256 processors of CITERRA. The mesh contains 6.2 million hexahedral elements and 1.2 billion degrees of freedom. The resolution at the surface is 60 m and the mesh is numerically accurate down to the maximum frequency content of the seismic source (6 Hz). The mesh is created in 5 min of user-time: 60% of this time is dedicated to geometry creation, and only 20% to evaluation at the GLL points and assignment of material properties. Figure 3 shows the mesh and synthetic seismograms with a duration of 20 s for the main seismic event of the 1983 caldera unrest.

In summary, we have developed a parallel, semiautomatic mesher for a general geological layer-cake structure, capable of accommodating surface topography and the main internal discontinuities. The user provides only the shapes of the interfaces and the desired numerical resolution (balanced against the time step, i.e., against the numerical cost). The mesher includes some Fortran routines taken from the SPECFEM3D_BASIN solver and which allow the execution of CUBIT in an MPI environment. The main limitation is due to the requirement that each slice has the exact same number of hexahedra. It is likely that in the next version of the solver this limitation will be overcome, opening new possibilities for simulating seismic wave propagation in complex geological volumes. The ultimate goal is to mesh sedimentary basins fairly automatically. Another considerable limitation of this meshing strategy is the use of Fortran as the language to create the journal files, an approach that does

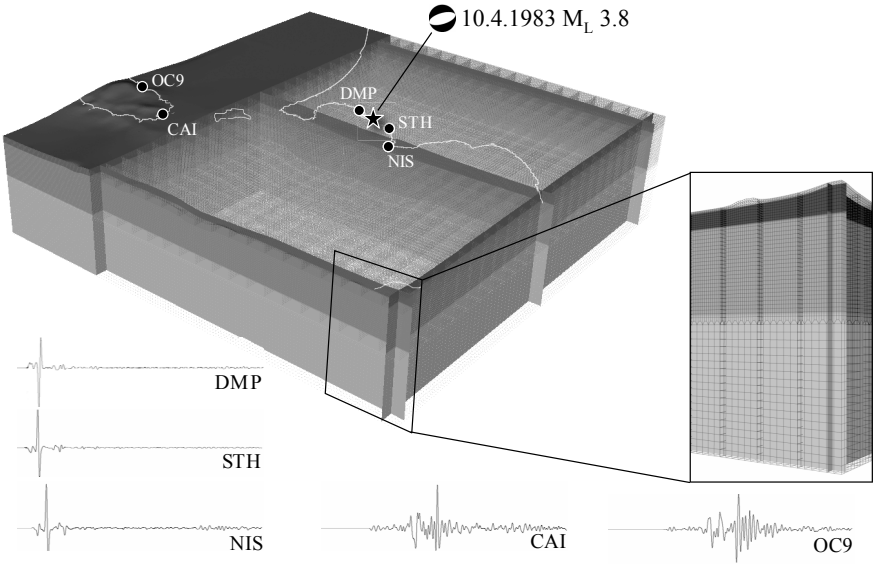

Fig. 3. CUBIT Mesh adopeted for the simulation on Campi Flegrei caldera. The seismograms show 20 s of ground motion associated with the main event of the 1982–84 crisis. The mesh is composed of more than six million elements and is accurate for the propagation of seismic waves with frequencies up to 6 Hz. The close up highlights the single refinement layer.

not provide the possibility to "interact" with the mesh during its creation. A significant step towards major flexibility is using the Python Interface, which is contained in the CUBIT package and provide the user with access to CUBIT commands from a python program. To take advantage of this versatility, we have developed some Python routines wrapping the fortran codes. In this case, the MPI environment is provided by PyMPI, a parallel interpreter of Python [16]. First, we have reproduced the same kinds of meshes for the layer-cake model class that we have presented in this subsection. The flexibility of CUBIT Python Interface enables us to handle fairly automatically more complex cases, such as sedimentatary basins.

3.2 Sedimentary basins

Our interest in modelling alluvial basins is related to the seismic properties of these regions. Since they involve high wave-speed contrasts between alluvial deposits and bedrock, seismic waves are trapped inside the basin, producing reverberations that are stronger and longer than in the surrounding areas, consequently involving higher seismic hazard.

In this section we introduce a serial CUBIT meshing strategy for basins. The technique is based upon two standard steps: 1) webcut the entire volume in small slices, and 2) mesh each slice with a standard scheme (e.g., the pave meshing scheme applied on one of the surfaces and swept along towards the opposing surface). Since the entire process is executed on one processor, we do not have to take extra care to check the conformity of the mesh. This technique allows us to decompose the basin into hexahedra, honoring the geometrical interfaces, but worsens the quality of the resulting mesh with respect to the strategy developed in the previous section.

As an example, we consider a mesh of the Grenoble valley (France) in the framework of a numerical benchmark of 3D ground motion simulations (EGS06, [22]). The webcutting decomposition is performed by creating a transition belt of slices around the outline of the basin, both externally and internally to the basin (Figure 4). These transition zones should accommodate any quality problems that arise after the refinement of the elements inside the basin. The resulting mesh is composed of 220 thousand elements with a resolution of 21 m inside the valley, and up to 900 m in the bedrock. The mesh is designed to propagate seismic waves accurately up to frequencies of approximately 3 Hz. The last step involves partitioning the global mesh in order to assign the corresponding part of the mesh to each processor for seismic wave propagation. The velocity snapshots shown in Figure 4 are produced by the spectral-element software package GeoELSE [21], which resolves the problem of partitioning and load balancing by incorporating the METIS software library [13]. This task is completely independent of the meshing process. The creation of the mesh requires detailed knowledge of CUBIT and its meshing schemes. In particular, the webcut decomposition is a long and tedious step that can take days. The entire process requires one or two weeks of dedicated time for an experienced user.

This strategy forms the basis of an automated way to handle mesh generation for a sedimentary basin in parallel. The logical scheme of the problem is depicted in Figure 5a. The procedure takes advantage of the Python Interface to have access to CUBIT. The strategy is based upon the idea of creating a so-called partitioner from a coarse mesh on the surface. This is the tool we will use to decompose the model and create the MPI communication map between the slices. To understand how the mesh is built in practice, we show the procedure for a conceptual model in Figure 5. Before proceeding with the actual meshing process, the user should prepare in a general XYZ ASCII format: 1) the topography (labeled S) and the interfaces of the layer-cake volume (one of these is a surface that accommodates the deeper parts of the basin), 2) the internal outline curve (labeled I) of the part of the basin that is honored, and 3) the outline curve (labeled B) of the basin, surrounded by 4) a transition curve (labeled T) (Figure 5a). Let us emphasize that the creation of the transition curve is the only intervention required from the user in this fairly automatic process. Furthermore, the user can add to the process some personal commands by means of a standard CUBIT journal file in order to adapt

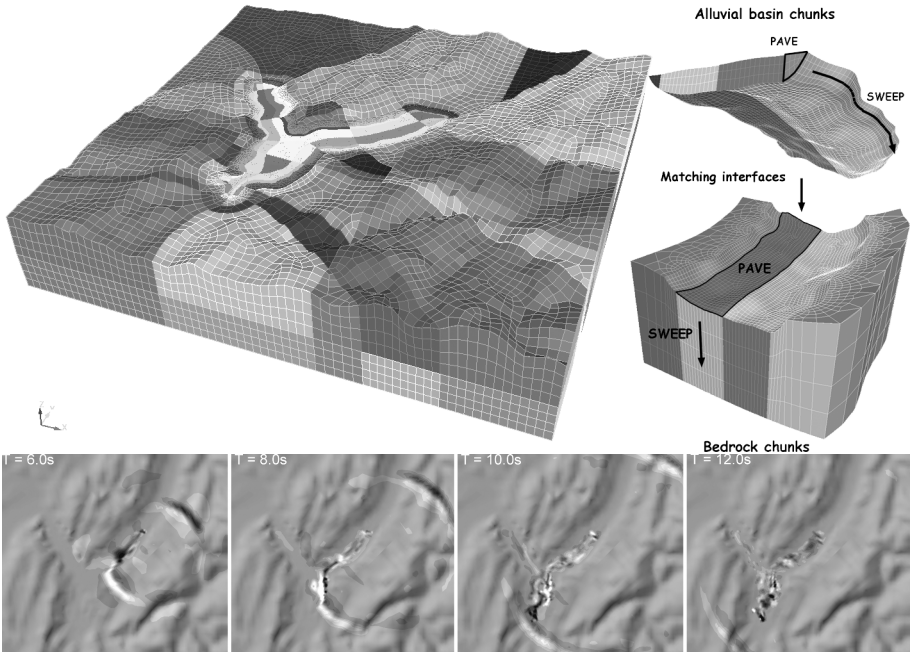

Fig. 4. (a) 3D numerical model used for the simulations of ESG06 "Grenoble Benchmark" with the GeoELSE code. The computational domain is subdivided into small chunks and each one is meshed starting from the alluvial basin down to the bedrock. (b) Snapshots of the velocity component for the Strong case 1, ESG06 "Grenoble Benchmark" [22].

the general case to a specific problem (e.g., the presence of a fault). Curves I and T are projected and imprinted on surface S, which is subsequently partitioned in three surfaces S_I, S_E, and S_T. The first two, S_I and S_E, are meshed with quadrilaterals using the pave scheme. The quad size is defined by the user and is related to the number of available processors. Each quadrilateral that shares at least one side with curves I and T is transformed into a curved element as in 5b. The group of these elements is the partitioner; one processor is assigned to each quad and the MPI map, with all the communication addresses, is created. In the MPI environment provided by PyMPI, each processor webcuts the volume V with the assigned element of the partitioner and keeps only its corresponding part, which it meshes independently following the procedure described for the layer-cake model (Figure 5c). In accordance with the MPI map, it sends (and receives) the position of the nodes on the lateral surface at the neighboring slices, so that a conforming mesh can be assured by the adjustment mentioned above. The hexahedral elements of the mesh that are inside the basin are refined in accordance with the constraints imposed by the wave speeds. Surface S_T is created in order to facilitate horizontal refine-

ments around the outline of the basin B. We construct this surface as a single slice without any partitioning so that the entire transition refinement sheet is built inside a single chunk (i.e., by a single processor). Consequently, we can avoid propagating the information of the refinement to some independent adjacent chunks, resulting in a higher quality mesh and a simpler MPI communication map. The resulting mesh honors the deep sections and the outline of the basin, concentrating horizontal refinement only where it is needed, thus saving a huge number of elements. Nevertheless, with respect to the serial procedure, the outcrop of the basin is generally not honored. This is possible only if the outcropping angle is fairly constant (with respect to the hexahedral dimensions of the simulation). On the contrary, if the outcropping angle changes along curve B (as in the case of the Taipei basin), a better solution is to use the spectral polynomial interpolation, since a robust fully 3D unstructured algorithm for hexahedra is still not available (the plastering meshing scheme that is currently under development [20] should be very useful in this context). Further study is required in order to understand when the interpolation is acceptable. The Python code that drives the meshing process is also applicable if more than one basin is present in the volume. Since each slice is different both in terms of the number of elements and in terms of the number of vertical side surfaces, the mesh is not well balanced. Therefore, before the simulation of seismic wave propagation METIS/ParMETIS should be used to partition the global graph, as is done in the GeoELSE package, thereby producing a load-balanced mesh and minimizing MPI communications.

4 Conclusion

In the coming Petaflops era, the SEM should become a standard tool for the study of seismic wave propagation, both for forward and inverse problems. The more the power provided by computer clusters, the higher the resolution that is available for the simulations. Consequently, the definition of a good geological model and the creation of an all-hexahedral unstructured mesh are critical. While a full 3D unstructured algorithm is still awaited, we have presented some meshing examples for typical seismological problems, such as layer-cake volumes and sedimentary basins. The adopted strategies are valid for classes of models that have the same logical geometrical representation. The procedures work automatically (or with minimal user intervention) embedded in a parallel MPI environment, thereby easily creating unstructured meshes with hundreds of millions of elements.

Acknowledgments

Emanuele Casarotti is supported by Project SP3D, Marie Curie Fellowship UE/INGV n. MOIF-CT-2005-008298. Marco Stupazzini is supported by SPICE, MRTN-CT-2003-504267. The numerical simulations for this research

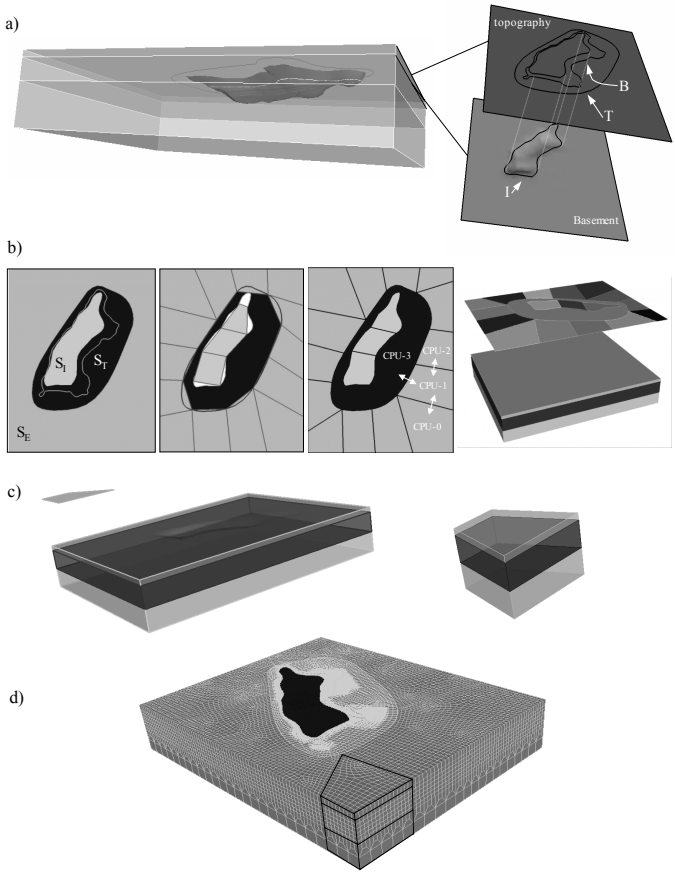

Fig. 5. Mesh of a conceptual model composed of a basin embedded in a layer-cake volume. For visual simplicity, we assume that all the surface are planar. (a) The "ingredients" that the user should provide are: interfaces, outline of the basin, internal outline corresponding to the part of the basin basement honored by the mesh, and a transition curve. (b) The creation of the "partitioner": from left to right, 1) the imprinting of the curves on the topography, 2) quad mesh of S_E and S_I, 3) transformation of the quads into curved elements (we note that S_T is a single slice), 4) the partitioner with respect of the mesh that should be meshed. (c) Each element of the partitioner is assigne to a processor and it is used as tool for the webcut. (d) The resulting mesh obtained reassembling the slices meshed by the independent processors; note the refinement in both the vertical and horizontal directions. The latter (lighter gray) is entirely included in the transition zone.

were performed on Caltechs Division of Geological & Planetary Sciences Dell cluster.

References

1. Bielak, J, Ghattas, O, and Kim, E J (2005) Parallel octree-based finite element method for large-scale earthquake ground motion simulation. Computer Modeling in Engineering and Sciences, 10(2):99–112.
2. Bouchon, M. and Sanchez-Sesma, F. (2007). Boundary integral equations and boundary elements methods in elastodynamics. In: Wu R-S and Maupin V (eds), Advances in Wave Propagation in Heterogeneous Media, Advances in Geophysics, 48 Elsevier - Academic Press.
3. `citerra.gps.caltech.edu`
4. Courant, R, Friedrichs, K O, and Lewy, H (1928) Über die partiellen differenzgleichungen der mathematischen physik (on the partial difference equations of mathematical physics). Mathematische Annalen, 100:32–74.
5. `cubit.sandia.gov`
6. Di Prisco C, Stupazzini M and Zambelli C (2007) Non-linear SEM numerical analyses of dry dense sand specimens under rapid and dynamic loading. International Journal for Numerical and Analytical Methods in Geomechanics, 31-6:757–788.
7. Faccioli E, Vanini M, Paolucci R and Stupazzini M (2005) Comment on "Domain reduction method for three-dimensional earthquake modeling in localized regions". Bulletin of the Seismological Society of America, 95:763–769
8. Käser M, Dumbser M, De la Puente J and Igel H (2007) An arbitrary high order discontinuous Galerkin method for elastic waves on unstructured meshes iii: Viscoelastic attenuation. Geophysical Journal International, 168:224–242
9. Komatitsch, D, Barnes, C, and Tromp, J (2000) Simulation of anisotropic wave propagation based upon a spectral element method. 65(4):1251–1260.
10. Komatitsch, D, Martin, R, Tromp, J, Taylor, M A, and Wingate, B A (2001) Wave propagation in 2-D elastic media using a spectral element method with triangles and quadrangles. Journal of Computational Acoustic, 9(2):703–718.
11. Komatitsch, D, Liu, Q, Tromp, J, Süss, P, Stidham, C, and Shaw, J H (2004) Simulations of ground motion in the Los Angeles basin based upon the Spectral-Element method. Bulletin Seismological Society American, 94:187–206.
12. Komatitsch, D, Tsuboi, S, and Tromp, J (2005) The spectral-element method in seismology. In: Levander A. and Nolet G., (eds), Seismic Earth: Array Analysis of Broadband Seismograms, volume 157 of Geophysical Monograph, pages 205–228 American Geophysical Union, Washington DC, USA.
13. `glaro.sdtcumn.edu/gkhome/views/metis`
14. Lee S J, Chen H W, Liu Q, Komatitsch D, Huang B S, and Tromp J (2007) Mesh generation and strong ground motion simulation in the taipei basin based upon the spectral-element method. submitted.
15. Moczo P, Robertsson J O A, and Eisner L (2007) The finite-difference time-domain method for modeling of seismic wave propagation, in advances in wave propagation in heterogeneous Earth In: Wu R-S and Maupin V (eds), Advances in Wave Propagation in Heterogeneous Media, Advances in Geophysics, 48 Elsevier - Academic Press.

16. `pympi.sourceforge.net`
17. Ronchi C, Ianoco R and Paolucci P S (1996) The "Cubed Sphere": a new method for the solution of partial differential equations in spherical geometry. J Comput Phys, 124:93–114.
18. Satriano C, Zollo A, Capuano P, Russo G, Vanorio T, Caielli G, Lovisa L and Moretti M (2006) A 3D velociy model for earthquake location in Campi Flegrei area application to the 1982-84 uplift event. In: Zollo A, Capuano P and Corciulo M (eds.), Geophysical Exploration of the Campi Flegrei (Southern Italy) caldera interiors: Data, Methods and Results. DoppiaVoce ed, Naples.
19. Scott M A, Earp M N, Benzley S E and Stephenson M B (2005) Adaptive sweeping techniques. In Proceedings of the 14th International Meshing Roundtable, pages 417–432 Springer.
20. M L Staten , R A Kerr, S.J.Owen and T D Blacker (2006). Unconstrained paving and plastering: Progress update. In: Proceedings, 15th International Meshing Roundtable, pages 469–486. Springer-Verlag.
21. Stupazzini M (2004) A spectral element approach for 3D dynamic soil-structure interaction problems. PhD thesis, Milan University of Technology, Milan, Italy.
22. Stupazzini M. (2006) 3D ground motion simulation of the grenoble valley by geoelse. In: Third International Symposium on the Effects of Surface Geology on Seismic Motion, August 30th - September 1st, Grenoble, France.