A Preconditioned 3-D Multi-Region Fast Multipole Solver for Seismic Wave Propagation in Complex Geometries

S. Chaillat[1,3,*], J. F. Semblat[2] and M. Bonnet[3]

[1] *Georgia Tech, College of Computing, Atlanta, USA.*
[2] *Université Paris Est, IFSTTAR, Paris, France.*
[3] *POems (CNRS-ENSTA-INRIA), Appl. Math. Dept., ENSTA, Paris, France.*

Abstract. The analysis of seismic wave propagation and amplification in complex geological structures requires efficient numerical methods. In this article, following up on recent studies devoted to the formulation, implementation and evaluation of 3-D single- and multi-region elastodynamic fast multipole boundary element methods (FM-BEMs), a simple preconditioning strategy is proposed. Its efficiency is demonstrated on both the single- and multi-region versions using benchmark examples (scattering of plane waves by canyons and basins). Finally, the preconditioned FM-BEM is applied to the scattering of plane seismic waves in an actual configuration (alpine basin of Grenoble, France), for which the high velocity contrast is seen to significantly affect the overall efficiency of the multi-region FM-BEM.

1 Introduction

Due to rapid and steady increase of available computational resources, the simulation of wave propagation in 3D configurations is currently a very active research area. The main advantage of the boundary element method (BEM) is that only the domain boundaries (and possibly interfaces) are discretized, leading to a reduction of the number of degrees of freedom (DOFs), and avoiding cumulative effects of grid dispersion [17, 18]. The BEM is well suited to unbounded-domain idealizations commonly used in seismology, as exact

---

*Corresponding author. *Email addresses:* `stephanie.chaillat@ensta-paristech.fr` (S. Chaillat), `semblat@lcpc.fr` (J. F. Semblat), `bonnet@lms.polytechnique.fr` (M. Bonnet)

satisfaction of radiation conditions is built into the formulation [3, 20]. However, the standard BEM leads to fully-populated matrices, which result in high computational costs in CPU time ($\mathcal{O}(N^2)$ per iteration using an iterative solver such as GMRES) and memory requirements ($\mathcal{O}(N^2)$), where N denotes the number of DOFs of the BEM model. The appearance of accelerated boundary element (BE) methodologies, allowing complexities far lower than those of traditional BEMs, has dramatically improved the capabilities of BEMs for many areas of application, largely owing to the rapid development of the Fast Multipole Method (FMM) over the last two decades [23]. Such approaches have resulted in considerable solution speedup, memory requirement reduction, and model size increase. The FMM is inherently associated with iterative solvers (usually GMRES), and is known to require $\mathcal{O}(N\log N)$ CPU time per iteration for Helmholtz-type equations [9, 10, 34]. To date, only few studies have been devoted to the FMM in elastodynamics (including [15] for the frequency-domain case, [35] for the time-domain case and [4] for a formulation specialized to surface waves), whereas FMMs for the Maxwell equations have been more extensively investigated, see e.g. [9, 16, 21, 34]. The present authors recently proposed an elastodynamic single-domain FM-BEM formulation which incorporates recent advances of FMM implementations for Maxwell equations [6], with BEM models of size up to $N = \mathcal{O}(10^6)$ run on a single-processor PC, then extended it to multi-domain situations, with emphasis on alluvial-basin configurations, by developing a FMM-based BE-BE coupling approach suitable for 3-D piecewise-homogeneous media [7].

The previous studies [6, 7] revealed that iteration count could significantly hinder the overall efficiency of the elastodynamic FM-BEM, especially for multi-region configurations, and even for problem sizes well within the computational platform's limitations in terms of required memory and CPU cost of a single iteration. This article aims at addressing this issue via a simple preconditioning approach based on an inner-outer GMRES algorithm, whose usefulness is then demonstrated on 3-D numerical examples representative of seismic wave propagation. The proposed preconditioned FM-BEM is then applied to a more realistic seismological configuration, namely the propagation of seismic waves in an alpine basin (Grenoble, France), for which the high velocity contrast is seen to significantly affect the overall performance of the multi-region FM-BEM. The paper is organized as follows. Classical concepts pertaining to elastodynamic BEM and FMM are summarized in Section 2. The preconditioning strategy is presented and demonstrated on numerical examples in Section 3, and applied on the Grenoble site model in Section 4.

2 Standard and fast multipole accelerated BEM

2.1 Single-region boundary element method

Let Ω denote a region of space occupied by an isotropic elastic solid characterized by μ (shear modulus), ν (Poisson's ratio) and ρ (mass density). A time-harmonic motion with circular frequency ω is assumed, and the implicit factor $e^{-i\omega t}$ will be, as usual, omitted throughout. Assuming the absence of body forces, the displacement $\boldsymbol{u}$ is given at an

interior point $x \in \Omega$ by the well-known integral representation formula:

$$u_k(x) = \int_{\partial \Omega} [t_i(y) U_i^k(x,y;\omega) - u_i(y) T_i^k(x,y;\omega)] \mathrm{d}S_y, \tag{2.1}$$

where t is the traction vector on the boundary $\partial \Omega$, and $U_i^k(x,y;\omega), T_i^k(x,y;\omega)$ are the i-th components of the elastodynamic fundamental displacement and traction, respectively, generated at $y \in \mathbb{R}^3$ by a unit point force applied at $x \in \mathbb{R}^3$ along the direction k [13]:

$$U_i^k(x,y;\omega) = \frac{1}{k_S^2 \mu} \left((\delta_{qs}\delta_{ik} - \delta_{qk}\delta_{is}) \frac{\partial}{\partial x_q} \frac{\partial}{\partial y_s} G(|y-x|;k_S) + \frac{\partial}{\partial x_i} \frac{\partial}{\partial y_k} G(|y-x|;k_P) \right),$$

$$T_i^k(x,y;\omega) = \mu \left[\frac{2\nu}{1-2\nu} \delta_{ij}\delta_{k\ell} + \delta_{ik}\delta_{j\ell} + \delta_{jk}\delta_{i\ell} \right] \frac{\partial}{\partial y_\ell} U_h^k(x,y;\omega) n_j(y), \tag{2.2a}$$

where

$$k_S^2 = \frac{\rho\omega^2}{\mu}, \quad k_P^2 = \frac{1-2\nu}{2(1-\nu)} k_S^2$$

in which $n(y)$ is the unit normal to $\partial \Omega$ directed outwards of Ω and $G(r;k)$, defined by

$$G(r;k) = \frac{\exp(ikr)}{4\pi r} \tag{2.3}$$

is the free-space Green's function for the Helmholtz equation with wavenumber k corresponding to either P or S elastic waves.

The boundary traces $(u,t)|_{\partial \Omega}$ featured in the integral representation (2.1) are governed by the singular boundary integral equation:

$$c_{ik}(x) u_i(x) + \fint_{\partial \Omega} u_i(y) T_i^k(x,y;\omega) \mathrm{d}S_y - \int_{\partial \Omega} t_i(y) U_i^k(x,y;\omega) \mathrm{d}S_y = 0, \tag{2.4}$$

where $\fint$ indicates a Cauchy principal value (CPV) singular integral and the *free-term* $c_{ik}(x)$ is equal to $0.5\delta_{ik}$ in the usual case where $\partial \Omega$ is smooth at x.

To formulate a well-posed problem from (2.4), part of the boundary traces $(u,t)|_{\partial \Omega}$ must be prescribed by applying boundary conditions in the usual way, and ω must not coincide with one of the countable (real or fictitious) eigenfrequencies of the integral operator governing the part of $(u,t)|_{\partial \Omega}$ that remains unknown. The numerical solution of (2.4) is then based on a boundary element (BE) discretization of the surface $\partial \Omega$ and boundary traces (u,t), here based on linear three-noded triangular boundary elements, piecewise-linear continuous (i.e. isoparametric) interpolation of displacements and a piecewise-constant interpolation of tractions. This leads to the system [3]:

$$HU + GT = 0, \tag{2.5}$$

where H,G are fully populated, nonsymmetric, influence matrices while vectors U,T, collect the displacement and traction degrees of freedom (DOFs). Moreover, a collocation approach is employed here whereby rows of (2.5) are generated by enforcing (2.4)

at a discrete set of points x. Upon introduction of boundary conditions, the matrix equation (2.5) is recast in the form:

$$Kx = F, \qquad (2.6)$$

where the N-vector x collects the unknown DOFs, while the $N \times N$ matrix of influence coefficients K contains the columns of H and G associated with the unknown DOFs.

BEM matrix equations such as (2.6) are here solved iteratively using the GMRES algorithm [29]. The influence matrix K being fully-populated, standard BEM numerical integration procedures entail a $\mathcal{O}(N^2)$ computing time for each GMRES iteration through the evaluation of Krylov vectors $x^{k+1} = Kx^k$. To lower this high complexity, which is unacceptable for large BEM models, fast BEM solutions techniques such as the Fast Multipole Method (FMM) used here are called for.

2.2   Fast Multipole Method: principle

The goal of the FMM is to speed up the matrix-vector product computation required for each iteration of the iterative solver applied to the BEM-discretized equation (2.6) while avoiding actual computation and storage of the BEM influence matrix. Substantial savings in both CPU time and memory are thus achieved.

In general terms, the FMM exploits a reformulation of the fundamental solutions in terms of products of functions of x and of y, so that (unlike in the traditional BEM) integrations with respect to y can be reused when the collocation point x is changed. On decomposing the position vector $r = y - x$ into $r = (y - y_0) + r_0 - (x - x_0)$, where x_0 and y_0 are two poles and $r_0 = y_0 - x_0$ and invoking the Gegenbauer addition theorem, the Helmholtz Green's function is written as [9, 12]:

$$G(|r|, k) = \lim_{L \to +\infty} \int_{\hat{s} \in S} e^{ik\hat{s}.(y - y_0)} \mathcal{G}_L(\hat{s}; r_0; k) e^{-ik\hat{s}.(x - x_0)} \, d\hat{s}, \qquad (2.7)$$

where S is the unit sphere of $\mathbb{R}^3$ and the *transfer function* $\mathcal{G}_L(\hat{s}; r_0; k)$ is defined in terms of the Legendre polynomials P_p and the spherical Hankel functions of the first kind $h_p^{(1)}$ by:

$$\mathcal{G}_L(\hat{s}; r_0; k) = \frac{ik}{16\pi^2} \sum_{0 \leq p \leq L} (2p+1) i^p h_p^{(1)} (k|r_0|) P_p \big(\cos(\hat{s}, r_0) \big). \qquad (2.8)$$

Representations of the form (2.7) with $\mathcal{G}_L$ replaced with suitably-defined (tensor) transfer functions are then easily found [6] for the elastodynamic fundamental solution (2.2).

Single-region FMM.   A 3D cubic grid of linear spacing d embedding the whole boundary $\partial\Omega$ is then introduced. The FMM basically consists in using decomposition (2.7), with the poles x_0 and y_0 being chosen as the cell centers, whenever x and y belong to *non-adjacent* cubic cells. The treatment of such "FM" contributions exploits the multipole

expansions of the fundamental solutions (2.2), truncated at a finite L. When x and y belong to adjacent cells, traditional BEM evaluation methods based on expressions (2.2) and (2.3) are used. As a result, the system matrix is split into two submatrices

$$K = K^{near} + K^{FM}.$$

The matrix K^{near} is explicitly assembled and stored, while K^{FM} is not. To improve further the computational efficiency of the FM-BEM, a multi-level scheme is used whereby cells are recursively subdivided into eight smaller cubic cells [9, 16, 21]. The cell-subdivision approach is systematized by means of an oct-tree structure of cells. The level $\ell = 0$, composed of only one cubic cell containing the whole surface $\partial \Omega$, is the tree root. The subdivision process is further repeated until the finest level $\ell = \bar{\ell}$, implicitly defined by a preset subdivision-stopping criterion ($d^{\bar{\ell}} \geq d^{min}$), is reached, where d^{ℓ} is the linear level-ℓ cell size. Level-$\bar{\ell}$ cells are usually termed leaf cells. Decompositions of type (2.7) diverge in the low-frequency limit $\omega \to 0$ [10]. FM-BEMs exploiting them are thus mid-frequency FMMs and, as established in various empirical studies (e.g. in [6, 9]), require leaf cells to have a minimum size relative to wavelength. Here, the bound $d^{min} = 0.30\lambda_S$ is used.

The theoretical complexity of the multi-level FMM is $\mathcal{O}(N\log N)$ per GMRES iteration both for CPU time and memory (see [6] for further details on the method and its implementation for single-domain elastodynamic problems).

Multi-region FMM.   This method is then extended to 3-D elastic wave propagation in piecewise homogeneous domains in the form of a FM-accelerated multi-region BE-BE coupling approach. Such formulation is useful e.g. for studying site effects (i.e seismic wave amplification) in sedimentary basins. The coupling approach rests upon applying the FMM independently for each homogeneous subregion. The BE-BE coupling then consists in defining suitable linear combinations of equations arising from each subregion, with weighting coefficients adjusted on the basis of numerical experiments. Further details on the BE-BE coupling and its implementation are available in [7].

3   Preconditioning strategy

No preconditioning was used in the previous studies [6, 7], which revealed that large numbers of GMRES iterations could constitute a significant limiting factor, especially for multi-region configurations, and even for problem sizes well below the expected limitations in memory or CPU cost of a single iteration. A preconditioning strategy is clearly needed to improve convergence properties for large coupled BE-BE models.

3.1   Overview of preconditioning strategies in the context of the FMM

Preconditioning strategies in the context of the FMM.   Generally speaking, preconditioning the linear system (2.6) consists in reformulating it using a preconditioning matrix

M which approximates *K* in some sense while being easier to evaluate (either explicitly or implicitly through matrix-vector products), store and invert. A right preconditioning strategy then consists of considering the system

$$KM^{-1}z = F, \quad \text{with} \quad Mx = z, \tag{3.1}$$

where *M* should lower the condition number of the system, i.e. $\text{Cond}(KM^{-1}) < \text{Cond}(K)$. Of course, forming the preconditioned matrix KM^{-1} is never needed in practice, linear systems with matrix *M* being solved instead.

The determination of an efficient preconditioner for the elastodynamic FM-BEM is a largely open issue. In [15], where problem sizes of at most $N = \mathcal{O}(10^4)$ are considered, a block-diagonal preconditioner is used. Preconditioning of electromagnetic FMMs is more developed, with several proposed approaches available in the literature. The incomplete LU factorization with threshold has been successfully applied in e.g. [33] to various electromagnetic scattering problems, in conjunction with the FMM. Other strategies exploit the sparse approximate inverse (SPAI) of *K*, defined as the matrix *M* minimizing $\|I - MK\|_F$ subject to sparsity constraints [1]. The flexible GMRES (FGMRES [27]), where the preconditioner may vary from one GMRES step to the next, has also been proposed for FMM implementations. In [5], the standard (accurate) FMM is applied to the matrix-vector products for the outer solver whereas the inner solver exploits a low-accuracy FMM, itself preconditioned using the SPAI. This method is shown to be efficient for problems of size up to $N = \mathcal{O}(10^6)$ DOFs.

3.2 Inner-outer preconditioning using the near influence matrix

In the FMM, the complete influence matrix *K* is never explicitly assembled, and K^{near}, a sparse matrix into which the near contributions are assembled, is the only naturally available matrix. It is proposed here to use K^{near} as the preconditioning matrix, i.e. to set $M = K^{\text{near}}$, in a right-preconditioning approach. The size of K^{near}, while permitting its assembly and storage (as a sparse matrix), is however too large (see the numerical examples thereafter) for applying direct solvers to it. For this reason, two GMRES solvers are nested in an inner-outer scheme, where the outer GMRES (restarted every *m* iterations) solves the linear system of interest while the inner GMRES (without restart) solves preconditioning linear systems based on K^{near}. Moreover, to avoid having to apply M^{-1} to the Krylov vectors as in the standard version of the right-preconditioned GMRES, recourse is made instead to the FGMRES variant [27] of the algorithm, where such inversion is replaced with storage of *m* additional auxiliary vectors. Even though FGMRES allows one to vary *M* at each step, the same *M* is used here for each step $k = 1, \cdots, m$. This leads to Algorithm 3.1, which has been implemented using the FGMRES routine `zPackfgmres.f` [14] and is demonstrated thereafter. The advantage of this preconditioning strategy is that the computation of the preconditioner is not CPU-consuming since the sparse matrix K^{near} is computed and stored anyway. The matrix-vector product needed for the inner GMRES solver takes advantage of the structure of the computation of the near contributions.

Algorithm 3.1: GMRES(m) with right preconditioning and inner GMRES for solving $Ax=b$.

$r_0=b-Ax_0$, $\beta=\|r_0\|_2$, $v_1=r_0/\beta$
while $\|r_0\| \geq \varepsilon_{\text{outer}}(\|A\|\|x_0\|+\|b\|)$ **do**
    for $j=1,\cdots,m$ **do**
        $Mz_j=v_j$ solved using inner GMRES (unpreconditioned, no restart, tolerance $\varepsilon_{\text{inner}}$)
        $w=Az_j$
        for $i=1,\cdots,j$ **do**
            $h_{i,j}=(w,v_i)$
            $w=w-h_{i,j}v_i$
        end for
        $h_{j+1,j}=\|w\|_2$, $v_{j+1}=w/h_{j+1,j}$
        $Z_m=[z_1,\cdots,z_m]$, $\bar{H}_m=\{h_{i,j}\}_{1\leq i\leq j+1,1\leq j\leq m}$
    end for
    $y_m=\operatorname{argmin}_y\|\beta e_1-\bar{H}_m y\|_2$, $x_m=x_0+Z_m y_m$
    $x_0=x_m$
    $r_0=b-Ax_0$, $\beta=\|r_0\|_2$, $v_1=r_0/\beta$
end while

3.3 Numerical validation

The efficiency of this preconditioning strategy is now checked on two variants of a seismo-logy-oriented problem often used as benchmark [7, 22, 30], where the scattering by either a semispherical canyon (case (a)) or a semispherical alluvial basin (case (b)) of plane seis-mic waves travelling in an elastic half-space is considered (Fig. 1). In both cases, an oblique ($\theta=30°$) incident plane P-wave is considered, the material in Ω_1 has a Poisson ratio $\nu^{(1)}=0.25$, the free surface is meshed within a disk of radius $b=5a$ (where a is the radius of the semispherical obstacle), and the non-dimensional frequency is set to $k_{\text{P}}^{(1)}a/\pi=2$.

In case (a) (Fig. 1(a)), the BE mesh features $N=111{,}237$ DOFs. In case (b), the alluvial basin Ω_2 (Fig. 1(b)) is also modeled (with material parameters in Ω_2 given by $\mu^{(2)}=0.3\mu^{(1)}$, $\rho^{(2)}=0.6\rho^{(1)}$ and $\nu^{(2)}=0.3$), and the mesh features $N=190{,}299$ DOFs.

Efficiency of the preconditioning strategy. All examples presented in this article have been run on the same 8-processor PC (RAM: 64 GB, CPU frequency: 2.33 GHz), with each FMM analysis performed independently on a single processor.

The number of iterations and total CPU time required by the complete solution pro-cedure for cases (a) and (b), with and without the preconditioning method of Section 3.2, are given in Table 1. The GMRES tolerances are set at $\varepsilon_{\text{inner}}=10^{-1}$ for the inner solver and $\varepsilon_{\text{outer}}=10^{-3}$ for the outer solver. No restart is used for the inner solver, while the outer solver is restarted every $m=50$ iterations. Preconditioning is seen to considerably reduce the outer iteration count. However, the inner GMRES iterations in the preconditioned version also contribute to the overall computational cost (noting, however, that inner
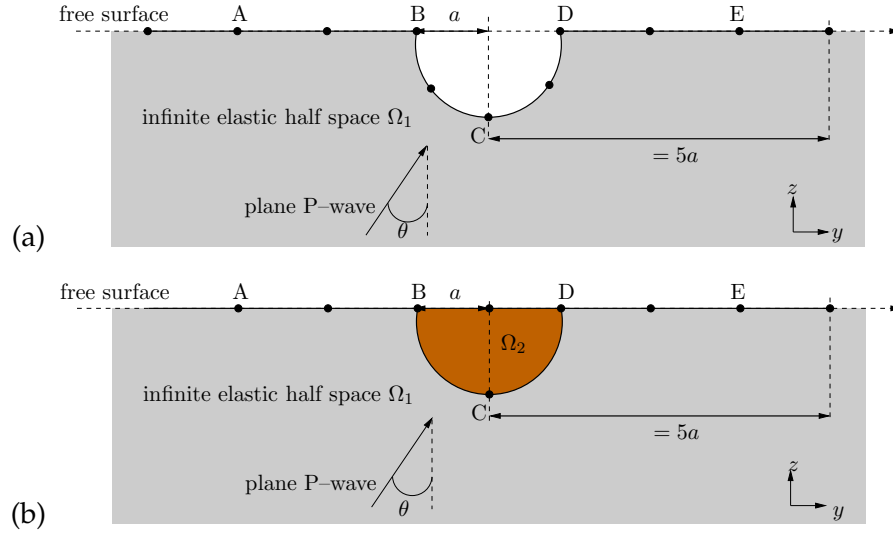
Figure 1: Diffraction of an oblique incident plane P-wave by a semi-spherical canyon (top) or a semi-spherical basin (bottom): geometry and notation.

Table 1: Diffraction of an incident plane wave by a semi-spherical canyon and basin: iteration counts and CPU time (with and without preconditioning).

Configuration	without prec.		with prec.		
	Iters	CPU	Iters (inner)	Iters (outer)	CPU
(a) canyon	43	33′19″	70	17	25′36″
(b) basin	388	7h59′27″	231	26	2h30′54″

matrix-vector products Mz are less expensive than outer products Kx). Thus, in spite of the large number of inner GMRES iterations performed by the inner-outer scheme, Table 1 shows the overall efficiency of the latter (measured in terms of cumulative CPU time) to be about three times that of the un-preconditioned version for case (b), which was known beforehand to be the more problematic one in terms of convergence speed.

Influence of the inner-solver tolerance parameter. An important parameter in our inner-outer scheme is the tolerance $\varepsilon_{\text{inner}}$ used in the stopping criterion of the inner GM-RES. To study its influence, $\varepsilon_{\text{inner}}$ is varied from $5\,10^{-2}$ to $5\,10^{-1}$, and the corresponding recorded values of total CPU time and iteration count are given in Table 2. Decreasing $\varepsilon_{\text{inner}}$ is seen to reduce the iteration count of the inner GMRES while increasing that of the outer GMRES. Since the latter iterations are more CPU-consuming than the former, efficient preconditioning requires finding a good compromise between the number of outer and inner iterations, which is achieved on this example by setting $\varepsilon_{\text{inner}} = 10^{-1}$. While the optimal choice of $\varepsilon_{\text{inner}}$ depends on the problem (through the ratio of the CPU times per inner and outer iterations), the value $\varepsilon_{\text{inner}} = 10^{-1}$ seems to be a reasonable default choice.

Table 2: Diffraction of an incident plane wave by a semi-spherical basin: influence of the tolerance $\varepsilon_{\text{inner}}$ used for the convergence of the inner solver.

$\varepsilon_{\text{inner}}$	nb iter. (inner)	nb iter. (outer)	CPU time
$5\ 10^{-2}$	338	25	3h12′34″
$8\ 10^{-2}$	248	25	2h35′43″
$1\ 10^{-1}$	231	26	2h30′54″
$3\ 10^{-1}$	164	41	2h50′44″
$5\ 10^{-1}$	171	58	3h31′43″

4   Modelling of an Alpine valley: Grenoble

The proposed preconditioned elastodynamic FMM is now applied to a more realistic seismological configuration, namely the diffraction of a vertical incident plane seismic wave by an Alpine valley (Grenoble, France).

4.1   Choice of the Grenoble site

The geological configuration, basin geometry and edges modify an incident wave field and may lead to large amplifications and higher signal duration. The Grenoble alpine valley (France), having previously been the subject of a numerical benchmark, offers an interesting case study with available mechanical and topographical data. This benchmark proposed various configurations (1D, 2D, 3D, with/without topography) to investigate weak (i.e linear) as well as strong (i.e. nonlinear) seismic motion in the Grenoble valley. Various numerical methods were used by different teams around the world, including classical BEM simulations performed in the 2D case [11]. The results are discussed in [2].

Geometry definition and mesh generation.   Mesh generation is a significant issue when dealing with seismological applications. For this preliminary study of a realistic site, the topography of the valley outside the sedimentary basin is not modeled to keep the BEM model size within manageable limits. The bedrock/sediment interface, known from the inversion of gravimetric anomalies (Vallon [36]), is defined through the topographical coordinates of the bedrock/sediment interface on a regular grid (every 250m).

   The horizontal geometry of the Alpine valley is depicted in Fig. 2 (left). The seismic recordings (accelerations) displayed in the figure show that the ground motion is strongly amplified in the basin (seismological stations *OGCU*, *OGSR*, *OGDH* and *OGPC*) when compared to the bedrock motion (station *OGMU*). To design the BEM model, the valley, which is Y-shaped when seen from above, is enclosed in a circle of radius $a \simeq 11.7$ km (Fig. 2, right). For this study, the meshed surrounding portion of the free surface is circumscribed within a disk of radius $D=30$ km ($\simeq 3a$). No topographical data are available at Γ_{12}^{a} and Γ_{12}^{b}. Consequently, the North ends of the Y-shaped valley are closed artificially, although the steep slopes thus introduced may induce artificial reflections at the
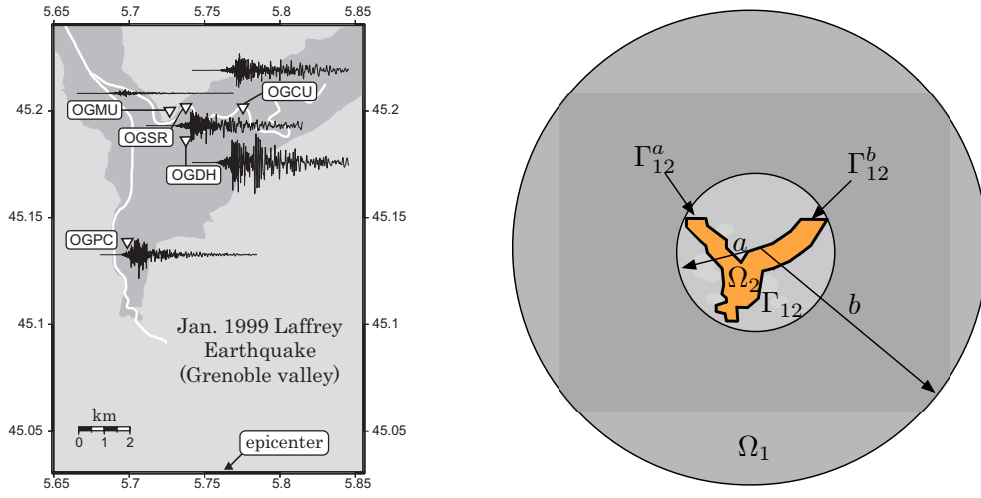
Figure 2: The Alpine valley: general topography and seismic motion recordings, from [32] (left); geometry and notation for the model (right).

basin edges. A study comparing various artificial valley closures should be done to evaluate their effects on the simulated seismic motion. However, for now this simple closure method is used in this study.

Material parameters. The bedrock and sedimentary basin models are as proposed in [8]. In the bedrock Ω_1 (Fig. 2), the P- and S-velocities and mass density are set to constant values:

$$c_P^{(1)} = 5,600 \text{ m.s}^{-1}, \quad c_S^{(1)} = 3,200 \text{ m.s}^{-1} \text{ and } \rho^{(1)} = 2,720 \text{ kg.m}^{-3}.$$

In the sedimentary basin, the velocity profile increases with depth z. The models used in numerical benchmark (ESG 2006) were: $c_P(z) = 1450 + 1.2z$ m.s^{-1}, $c_S(z) = 300 + 19\sqrt{z}$ m.s^{-1} and $\rho(z) = 2140 + 0.125z$ kg.m^{-3}, with z in meters. In this work, only a single homogeneous layer Ω_2 is used, with mean mechanical parameters set to:

$$c_P^{(2)} = 1,988 \text{ m.s}^{-1}, \quad c_S^{(2)} = 526 \text{ m.s}^{-1} \text{ and } \rho^{(2)} = 2,206 \text{ kg.m}^{-3}.$$

In [11], a piecewise homogeneous model of this deposit (involving 6 layers) was considered in the framework of a classical 2D BEM, which thus required a discretization of the interfaces between each subdomain. The results were satisfactory when compared to other numerical approaches accounting for the variations of the mechanical properties [8]. In another case study [31], the actual soil layering of the Volvi basin was modelled using different piecewise homogeneous BEM models (two to six layers), and detailed description of the geological structure was found to affect the propagation features only in the higher frequency range. The present preliminary model may thus be considered as reasonable for the low-frequency range, close to the fundamental frequency of the basin for which the most severe amplifications are known to occur [8].

Table 3: Propagation of an incident plane P-wave in an Alpine valley: computational data.

f	N	$\bar{\ell}_1;\bar{\ell}_2$	CPU time / iter (s)	iters (with prec.)	total CPU time
0.3 Hz	95,142	4;5	86.6	253	39h55′31″
0.6 Hz	141,288	5;6	77	747	75h45′44″

4.2 Surface displacements for a vertical incident plane P-wave

The diffraction by the valley of a vertical incident plane P-wave is considered for two frequencies: $f = 0.3$ Hz and $f = 0.6$ Hz. In Table 3, the number of DOFs N and the leaf levels $\bar{\ell}_1$ and $\bar{\ell}_2$ are given for the two frequencies together with the CPU time per iteration (without preconditioning), the number of iterations and the cumulative CPU time (with preconditioning). In Fig. 3, the modulus of the x-, y-, z-surface displacement components (normalized by the incident displacement amplitude) are displayed for the two frequencies. This computation shows the possibility of very high amplifications inside the alluvial basin (up to about 15.5 for $f = 0.6$ Hz). Most of the amplification occurs in the north part of the basin, for the z-component.

Limitations of the present FM-BEM for realistic seismic applications.   The BE meshes are designed so as to feature about 10 nodes per S-wavelength (Fig. 4), a common "rule of thumb" previously shown to be valid for our FM-BEM formulation [6,7]. Due to the high velocity contrasts between the bedrock and sediment (with $c_S^{(1)}/c_S^{(2)} \approx 6.1$), mesh conformity on Γ_{12} lead to mesh densities of about 10 nodes per *smallest* S-wavelength near Γ_{12}, i.e about 6 times too dense in the bedrock material surrounding Γ_{12} (Fig. 5).

   This arrangement leads to sub-optimal performance of the present FM-BEM. Standard complexity estimates for memory and CPU time are based on the assumption that the number of DOFs per wavelength is roughly uniform, resulting in roughly equal numbers of DOFs per leaf cell (due to the leaf cell size lower bound $d^{min} \geq 0.30\lambda_S$ empirically established in [6]). Here, the combined effect of highly heterogeneous mesh densities and cell size threshold leads to abnormally large numbers of DOFs in leaf cells intersecting Γ_{12}. This in turn degrades the sparsity of K^{near} and increases the memory required to hold it. The other consequence is that the CPU time and memory requirements are very sensitive to the number of levels. For the two frequencies studied $f=0.3$ Hz and $f=0.6$ Hz, the size of K^{near} is about 25 GB and 20 GB, respectively. The first remark is that K^{near} requires more memory for $f=0.3$ Hz than for $f=0.6$ Hz even though N is smaller; this is because the number of levels is larger in the latter case. Using the mesh with $N=141,288$ for $f=0.5$ Hz instead of $f=0.6$ Hz would lead to leaf levels $\bar{\ell}_1=4$ and $\bar{\ell}_2=5$ and result in an even larger K^{near} filling 52 GB. A similar effect makes the proportion of near contributions larger for $f=0.3$ Hz than for $f=0.6$ Hz, which explains the larger CPU time per iteration observed for $f=0.3$ Hz. So, although being based on relatively low BEM model sizes N (with limitations on the range of accessible frequencies) compared to the more academic examples treated in [6,7] with the same method (without preconditioning), the high ve-
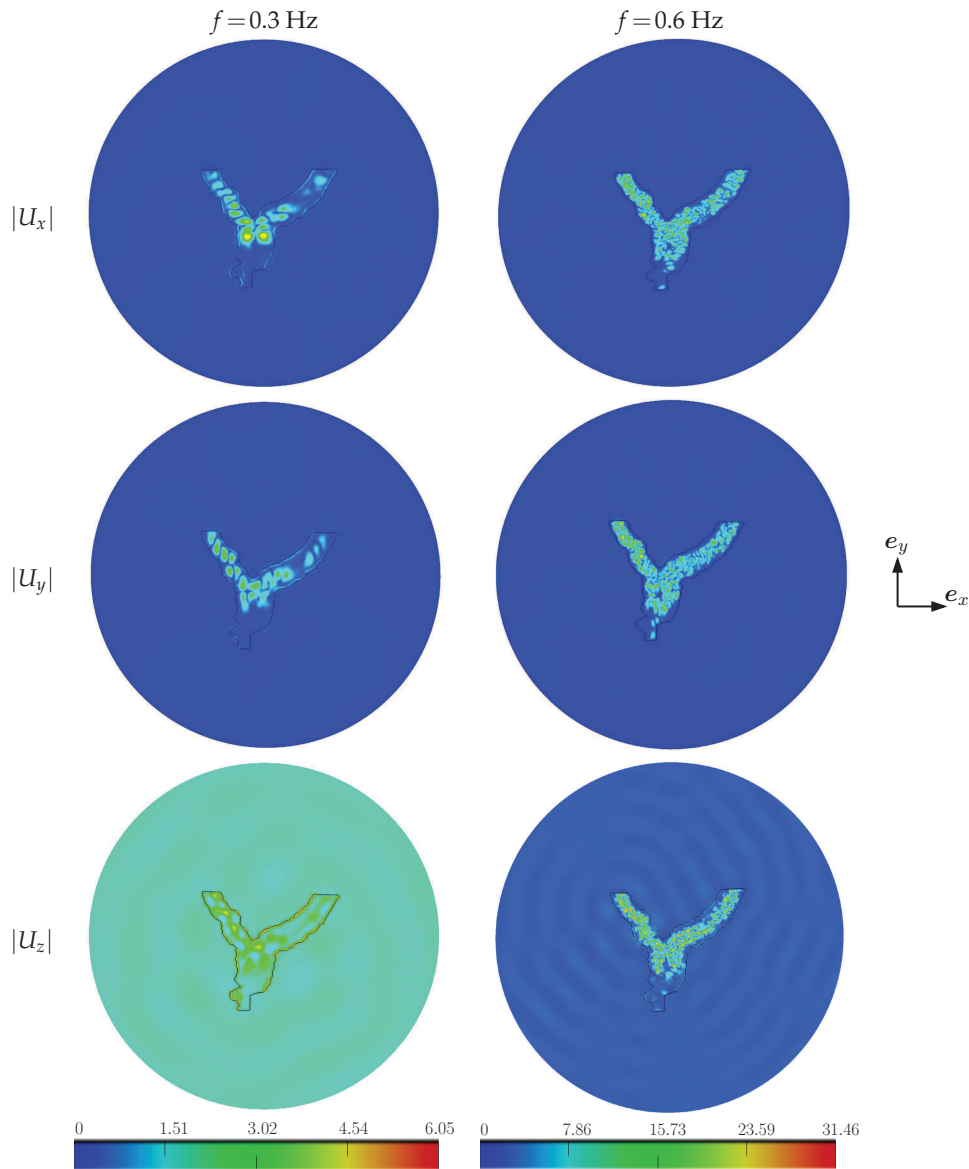
Figure 3: Propagation of a vertical incident plane P-wave in the Alpine valley: modulus of the x- (top), y- (middle) and z- (bottom) components of displacement for frequencies $f = 0.3$ Hz (left) and $f = 0.6$ Hz (right).

locity contrasts made computations performed for this study substantially more difficult. In fact, to enable the computation at $f = 0.6$ Hz, only about 8 points per S-wavelength are used in Ω_2 (instead of the 10 points per S-wavelength for $f = 0.3$ Hz).

The preconditioning strategy, while still useful, proved less efficient than in the examples of Section 3.3 because of the large size of K^{near}, increasing the cost of the inner
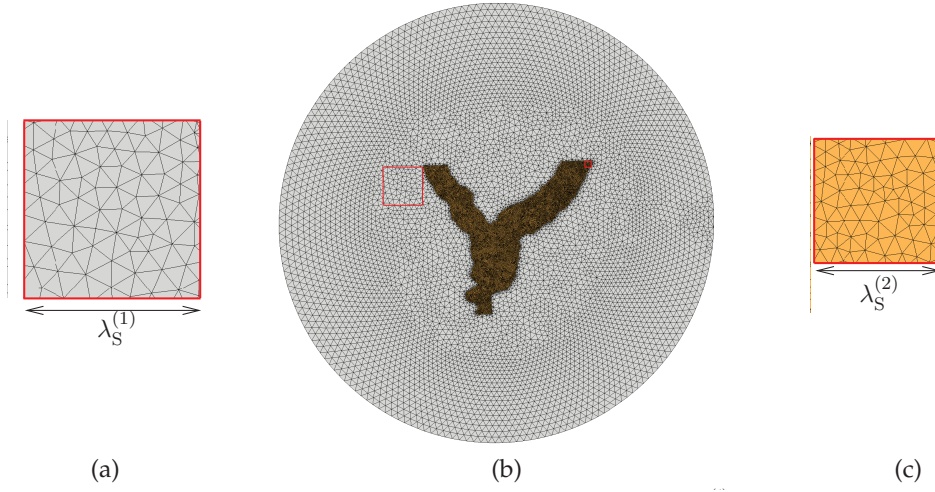
Figure 4: Density of points per S-wavelength (b) and close-ups on cells of size $\lambda_S^{(i)}$ (a and c), for $N=141,288$ and $f=0.6$ Hz.
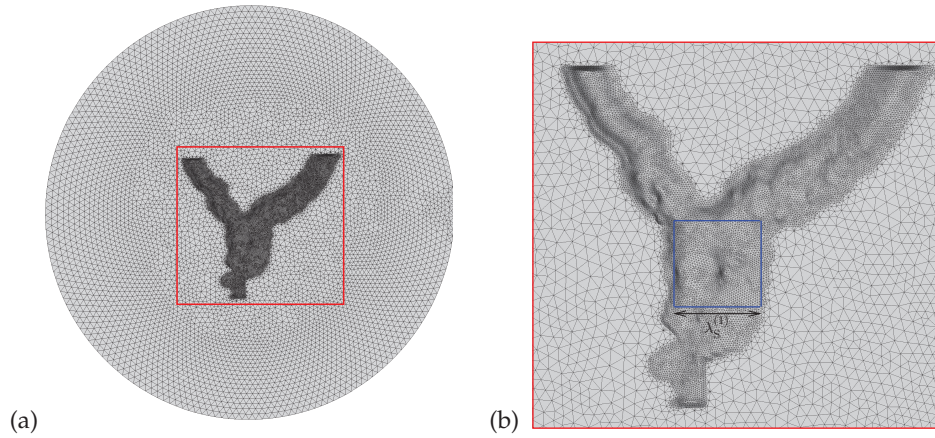

Figure 5: Density of points per S-wavelength in $\Omega^{(1)}$ (a) and close-up on the density of points per $\lambda_S^{(1)}$ on the interface Γ_{12} (b), for $N=141,288$ and $f=0.6$ Hz.

iterations. To mitigate this effect, the inner-solver tolerance was set to $\varepsilon_{inner}=5\,10^{-1}$ instead of the recommended value $\varepsilon_{inner}=10^{-1}$.

5   Conclusions and directions for future work

A simple and efficient preconditioning strategy has been proposed and implemented. This strategy is shown to be efficient on canyon problems and even more on basin prob-

lems (which are more ill-conditioned). It constitutes a step towards the definition of an optimal preconditioning strategy, a key issue for improving the efficiency of the elastodynamic FM-BEM for which further study is still needed. Moreover, a comparative study with the other preconditioning approaches used in electromagnetic FMM (incomplete LU, SPAI, inner-outer GMRES with two embedded FMM using various levels of accuracy), under way, is expected to bring worthwhile insight on this issue.

In the paper, the FM-BEM is also used to study a more realistic example: the diffraction of incident plane seismic waves by an Alpine valley (Grenoble, France). This example underlines the sub-optimal efficiency of the present FM-BEM for dealing with basin problems featuring a high velocity contrast between two layers, caused by non-uniform BE meshes near the interface. This issue would become even more problematic when considering different geological layers having various velocities.

To overcome this limitation, a method combining low- and mid-frequency FMM formulations could be used [19, 24]. Such treatment removes the $d^{\min} = 0.30\lambda_S$ lower bound for the linear size of leaf cells (the subdivision-stopping criterion used in the present mid-frequency FMM), allowing to instead define leaf cells in terms of a preset, uniform over the mesh, minimum number of DOFs per cell. Another possibility is to use non-conforming meshes [26]. Such approaches would permit the analysis of complex geological structures involving several layers having different wave velocities.

Acknowledgments

References

[1] G. Alléon, M. Benzi, L. Giraud, Sparse approximation inverse preconditioning for dense linear systems arising in computational electromagnetics, Numer. Algor., 16 (1997), 1–15.
[2] P. Y. Bard, E. Chaljub, C. Cornou, F. Cotton, and P. Gueguen, editors. Third Int. Symp. on the Effects of Surface Geology on Seismic Motion, Grenoble, France, 2006.
[3] Bonnet, M., Boundary Integral Equation Method for Solids and Fluids, Wiley, 1999.
[4] Ö Çakir, The multilevel fast multipole method for forward modelling the multiply scattered seismic surface waves, Geophys. J. Int., 167 (2006), 663–678.
[5] B. Carpentieri, I. S. Duff, L. Giraud and G. Sylvand, Combining fast multipole techniques and an approximate inverse preconditioner for large electromagnetism calculations, SIAM J. Sci. Comput., 27 (2005), 774–792.
[6] S. Chaillat, M. Bonnet and J.-F. Semblat, A multi-level fast multipole BEM for 3-D elastodynamics in the frequency domain, Comput. Meth. Appl. Mech. Engng., 197 (2008), 4233–4249.
[7] S. Chaillat, M. Bonnet and J.-F. Semblat, A new fast multi-domain BEM to model seismic wave propagation and amplification in 3D geological structures, Geophys. J. Int., 177 (2009), 509–531.

[8]   E. Chaljub, D. Komatitsch, J. P. Vilotte, Y. Capdeville, B. Valette and G. Festa, Spectral-element analysis in seismology. Adv. Geophys., 48 (2007), 365–419.

[9]   E. Darve, The fast multipole method: Numerical implementation, J. Comp. Phys., 160 (2000), 195–240.

[10] E. Darve and P. Havé, A fast multipole method for maxwell equations stable at all frequencies, Phil. Trans. Roy. Soc. (London), A362 (2004), 603–628.

[11] N. Delépine and J.F. Semblat. Site effects in a deep alpine valley for various seismic sources. Third Int. Symp. on the Effects of Surface Geology on Seismic Motion, Grenoble, France, 2006.

[12] M. A. Epton and B. Dembart, Multipole translation theory for the three-dimensional Laplace and Helmholtz equations, SIAM J. Sci. Comp., 16 (1995), 865–897.

[13] A. C. Eringen and E. S. Suhubi, Elastodynamics, II-linear theory, Academic Press, 1975.

[14] V. Frayssé, L. Giraud, and S. Gratton, A set of Flexible-GMRES routines for real and complex arithmetics. CERFACS Technical Report TR/PA/98/20, public domain software available on `www.cerfacs.fr/algor/Softs`, 1998.

[15] H. Fujiwara, The fast multipole method for solving integral equations of three-dimensional topography and basin problems, Geophys. J. Int., 140 (2000), 198–210.

[16] N. A. Gumerov and R. Duraiswami, Fast multipole methods for the Helmholtz equation in three dimensions, Elsevier, 2005.

[17] T. J. R. Hughes, A. Reali and G. Sangalli, Duality and Unified Analysis of Discrete Approximations in Structural Dynamics and Wave Propagation: Comparison of p-method Finite Elements with k-method NURBS, Comp. Meth. Appl. Mech. Engng., 197 (2008), 4104–4124.

[18] F. Ihlenburg and I. Babuška, Dispersion analysis and error estimate of Galerkin finite element methods for the Helmholtz equation, Int. J. Numer. Meth. Engng., 38 (1995), 3745–3774.

[19] L. J. Jiang and W. C. Chew, A mixed-form fast multipole algorithm, IEEE Trans. Antennas Propag., 53 (2005), 4145–4156.

[20] V. D. Kupradze, Dynamical problems in elasticity, Progress in solids mechanics (vol. 3), North Holland, 1963.

[21] Lu, C. C. and Chew, W. C., A multilevel algorithm for solving a boundary integral equation of wave scattering, Microwave Opt. Technol. Lett., 7 (1994), 466–470.

[22] T. K. Mossessian and M. Dravinski, Amplification of elastic waves by a three dimensional valley. Part 1: Steady state response. Earthquake Engng. Struct. Dyn., 19 (1990), 667–680.

[23] N. Nishimura, Fast multipole accelerated boundary integral equation methods, Appl. Mech. Rev., 55 (2002), 299–324.

[24] Y. Otani and N. Nishimura, A periodic FMM for Maxwell's equations in 3D and its application to problems related to photonic crystals. J. Comp. Phys., 227 (2008), 4630–4652.

[25] R. Y. S. Pak and B. B. Guzina, Seismic soil-structure interaction analysis by direct boundary element methods, Int. J. Solids Struct., 36 (1999), 4743–4766.

[26] T. Rüberg, Non-conforming Coupling of Finite and Boundary Element Methods in Time Domain, PhD thesis (TU Graz, Austria), 2007.

[27] Y. Saad, A flexible inner-outer preconditioned GMRES algorithm, SIAM J. Sci. Comput., 14 (1993), 461–469.

[28] Y. Saad, Iterative methods for sparse linear systems, SIAM, 2003.

[29] Y. Saad and M. H. Schultz, M. H., GMRES - a generalized minimal residual algorithm for solving nonsymmetric linear-systems, SIAM J. Sci. Statist. Comput., 7 (1986), 856–869.

[30] F. J. Sánchez-Sesma, Diffraction of elastic waves by three-dimensional surface irregularities,

Bull. Seism. Soc. Am., 73 (1983), 1621–1636.

[31] J. F. Semblat, M. Kham, E. Parara, P. Y. Bard, K. Pitilakis, K. Makra, D. Raptakis, Site effects: basin geometry vs soil layering, Soil Dynamics and Earthquake Eng. 25 (2005), 529–538.

[32] J. F. Semblat, A. Pecker, Waves and vibrations in soils: earthquakes, traffic, shocks, construction works, IUSS Press, 2009

[33] K. Sertel and J. L. Volakis, Incomplete LU preconditioner for FMM implementation, Microwave and Optical Technology Letters, 26 (2000), 265–267.

[34] G. Sylvand, La méthode multipôle rapide en éléctromagnétisme: performances, parallélisation, applications, Ph.D. thesis (ENPC, Paris, France), 2002, `http://pastel.paristech.org/308/`.

[35] T. Takahashi, N. Nishimura and S. Kobayashi, A fast BIEM for three-dimensional elastodynamics in time domain, Engng. Anal. Bound. Elem., 27 (2003), 491–506.

[36] M. Vallon. Estimation de l'épaisseur d'alluvions et sédiments quaternaires dans la région grenobloise par inversion des anomalies gravimétriques. Technical report, LGGE, Université Joseph Fourier, 1999. IPSN/CNRS (in French).