



HAL
open science

Fast Fourier homogenization for elastic wave propagation in complex media

Yann Capdeville, Ming Zhao, Paul Cupillard

► **To cite this version:**

Yann Capdeville, Ming Zhao, Paul Cupillard. Fast Fourier homogenization for elastic wave propagation in complex media. *Wave Motion*, 2015, 54, pp.170-186. 10.1016/j.wavemoti.2014.12.006 . hal-01275464

HAL Id: hal-01275464

<https://hal.science/hal-01275464v1>

Submitted on 21 Apr 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Fast Fourier homogenization for elastic wave propagation in complex media

Yann Capdeville^{a,*}, Ming Zhao^{a,b}, Paul Cupillard^{c,d}

^a*Laboratoire de Planétologie et Géodynamique de Nantes, CNRS, Université de Nantes, France*

^b*Laboratory of Computational Geodynamics, Chinese Academy of Sciences, Beijing, China*

^c*Équipe de sismologie, Institut de Physique du Globe de Paris, CNRS, France*

^d*GeoResources laboratory, Université de Lorraine, CNRS, Vandœuvre-lès-Nancy, France*

Abstract

In the context of acoustic or elastic wave propagation, the non-periodic asymptotic homogenization method allows one to determine a smooth effective medium and equations associated with the wave propagation in a given complex elastic or acoustic medium down to a given minimum wavelength. By smoothing all discontinuities and fine scales of the original medium, the homogenization technique considerably reduces meshing difficulties as well as the numerical cost associated with the wave equation solver, while producing the same waveform as for the original medium (up to the desired accuracy). Nevertheless, finding the effective medium requires one to solve the so-called “cell problem”, which corresponds to an elasto-static equation with a finite set of distinct loadings. For general elastic or acoustic media, the cell problem is a large problem that has to be solved on the whole do-

*Corresponding author

Email addresses: yann.capdeville@univ-nantes.fr (Yann Capdeville),
ming.zhao@univ-nantes.fr (Ming Zhao), paul.cupillard@univ-lorraine.fr
(Paul Cupillard)

main and its resolution implies the use of a finite element solver and a mesh of the fine scale medium. Even if solving the cell problem is simpler than solving the wave equation in the original medium (because it is time and source independent, based on simple tetrahedral meshes and embarrassingly parallel) it is still a challenge. In this work, we present an alternative method to the finite element approach for solving the cell problem. It is based on a well-known method designed by H. Moulinec and P. Suquet in 1998 in structural mechanics. This iterative technique relies on Green functions of a simple reference medium and extensively uses Fast Fourier Transforms. It is easy to implement, very efficient and relies on a simple regular gridding of the medium. Through examples we show that the method gives excellent results, even, under some conditions, for discontinuous media.

1. Introduction

Solving the elastic or acoustic wave equations in complex media is a difficult and a numerically expensive task, especially for media heterogeneous at scales much smaller than the minimum wavelength of the wavefield. For a given complex medium, the usual procedure to numerically model a wave propagation phenomena is first to mesh all the fine structures of the medium and then to solve the wave equation with our favorite solver. If the medium contains small scales, such a procedure is difficult and time consuming because, firstly, the mesh may be difficult to generate and secondly, the obtained fine and complex mesh induces a high numerical cost for the solver. An alternative to this simple but expensive approach is to pre-process the medium to compute an effective medium using an upscaling tool before meshing and

solving the wave equation. By smoothing out all the small scales from the medium, the upscaling step makes it possible to use a sparser and simpler mesh, leading to a lower numerical cost, for the wave equation solver. In many realistic situations, the medium presents no spatial periodicity, no natural scale separation or any kind of spatial statistical invariance. This difficulty excludes most of the classical and numerical homogenization techniques to upscale the medium. We use here the non-periodic homogenization technique [1, 2, 3], which is specifically designed to upscale such general deterministic media. If the non-periodic homogenization technique is strongly inspired from the classical two scale periodic homogenization [4], it has some strong differences as it will appear later on. One of them lies in the fact that the obtained effective properties are not spatially uniform, they are just “smoother” than the original medium.

One of the important research fields in which such general media are encountered is seismology. For many applications, seismologists work with limited frequency-band data of the ground motion recorded by seismic stations. This limited frequency band can be due to attenuation or instrument response but most of the time, it is simply the seismologist himself who limits the frequency content of his data using a band-pass filter. The reason to do so is linked to limited computing power resources available to model the data, but also to a limited knowledge of the Earth elastic structure. In the far-field of the source (an earthquake, for example), the fact that data has a maximum frequency f_{\max} ensures that the wavefield has a minimum wavelength λ_{\min} . Solving the seismic forward problem using numerical methods (such as finite differences, spectral elements, etc), that is solving the wave

equation to obtain the waveform at any space location, strongly relies on this knowledge of a λ_{\min} to accurately sample the wavefield. We assume that the elastic medium in which we need to solve the forward problem has a minimum size of characteristic heterogeneity λ_h . λ_h could be the shortest distance between two layers of a discontinuous medium or the fastest oscillation scale of a continuous medium. To estimate the scaling, as a function of λ_{\min} , of the computing time t_c necessary to solve the forward problem for a fixed signal duration, we need to distinguish two cases, depending on the regularity of the elastic medium under consideration:

1. if $\lambda_h \gg \lambda_{\min}$, we are in the smooth medium case (the wavefield oscillates much faster than the medium). In such a case, for N_s sources, the computing time t_c scales as

$$t_c \propto N_s \lambda_{\min}^{-(d+1)}, \quad (1)$$

where d is the problem dimension (2-D or 3-D). This is the optimal case in the sense that this scaling of t_c as a function of λ_{\min} can only be improved with some extra symmetries or assumptions on the medium.

2. if $\lambda_h \ll \lambda_{\min}$, we are in the rough medium case. In such a case, for N_s sources, the computing time t_c scales as

$$t_c \propto N_s \lambda_h^{-(d+1)}. \quad (2)$$

This second case is very common in most realistic applications. In practice, this $\lambda_h^{-(d+1)}$ scaling appears differently depending on the numerical solver used to solve the wave equation. For example, if finite elements are considered, then complex, fine and discontinuous structures lead to a complex

mesh which is usually difficult to generate and expensive to use. Indeed, in order to be accurate, the finite element mesh needs to honor all medium discontinuities. If finite differences are used, then small structures impose an expensive oversampling of the wavefield. The rough media case (case 2 above) is therefore a non-optimal configuration and a seismologist feels he is paying a computing price that he should not. This intuition is linked to the fact that it is well-known from observations that, somehow, waves of λ_{\min} wavelength are sensitive to small heterogeneity scales $\lambda_h \ll \lambda_{\min}$ only in an effective way and, if this effective medium was known, we could go back to the optimal scaling cost (case 1 above), that is a cost that scales with λ_{\min}^{-4} and not with λ_h^{-4} . This is exactly the objective of non-periodic homogenization [1, 2, 3]: finding the upscaling operator allowing us to compute the effective medium of a given rough medium so that the numerical cost scales as $t_c \propto N_s \lambda_{\min}^{-(d+1)}$ even if $\lambda_h \ll \lambda_{\min}$. The non-periodic homogenization method gets its name by opposition to the so-called two scale periodic homogenization [4] from which it is derived, a very powerful method but limited to periodic media. A sketch summarizing the non-periodic homogenization principle in the forward modeling context is shown in Fig. 1. The main idea of the method is to compute an effective version of the original medium for which meshing and computation are simpler and cheaper without degrading the waveform accuracy. It can be seen as a pre-processing step applied to the medium before importing it into the wave equation solver. It can also be seen as a generalization of the Backus averaging (or upscaling) technique [5]. Once the homogenized medium is obtained, any wave equation solver can be used, as long it can handle fully anisotropic and continuously varying media.

So far, we have justified the homogenization in the forward modeling context, but we could have done it for the inverse problem as well. Indeed, homogenization is very useful to build an inverse problem based on a multi-scale parameterization [6] and is already used to simplify the difficulties linked to the Earth heterogeneous crust [7, 8, 9] or to combine inverted models from different scales [10].

If the non-periodic homogenization leads to an optimal cost for the wave equation solver, the homogenization procedure itself requires one to solve the so-called cell problem, which can be an expensive step. This is in contrast to the classical periodic two scale homogenization whose cell problem is limited to one periodic cell of small size compared to the whole domain and is therefore inexpensive to solve, but is clearly limited to periodic media. For the non periodic homogenization method used here, the cell problem has to be solved on the whole domain. It can be solved globally (in such a case, a single cell including the whole elastic domain is used) or on multiple overlapping cells paving the domain. In any case, the cell problem implies a fine scale solver over the whole domain and is therefore a numerically expensive step (its computing cost scales as λ_h^{-d}). It is nevertheless time and source independent and can be embarrassingly parallelized (a program is “embarrassingly” parallel when different parallel tasks solving a global problem do not need communications between them to be completed). It is therefore a less expensive and complex problem than the original wave equation problem at the fine scale level. In previous works, the cell problem has been solved using a finite element solver [1, 2]. If finite element methods are essential

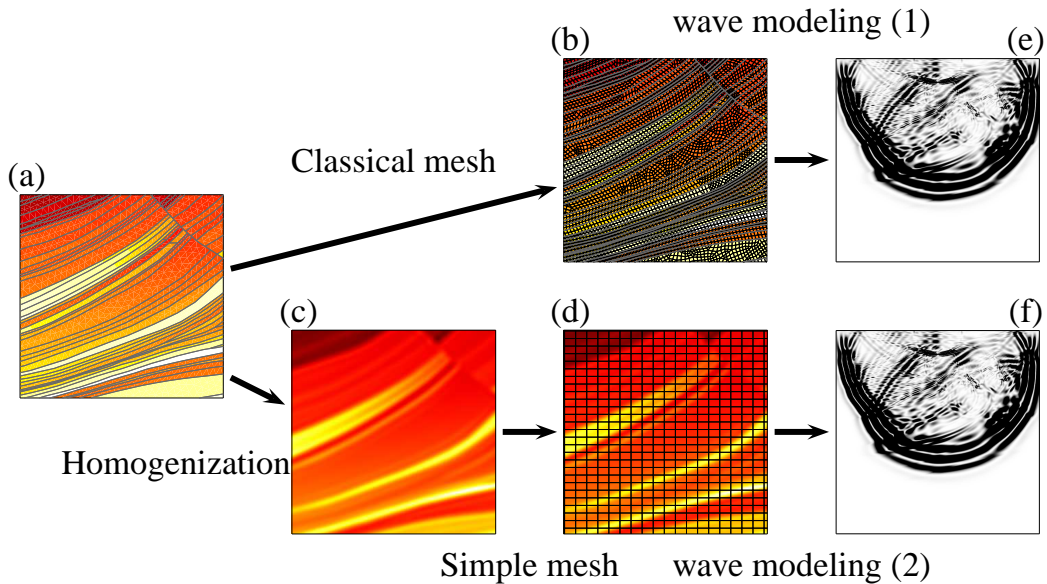


Figure 1: Homogenization principle. One wishes to propagate waves up to a maximum frequency f_{\max} in a complex elastic model (a) using a numerical solver (e.g. the spectral element method). The first and classical solution is to mesh all discontinuities of (a) to obtain the complex mesh (b) (here a spectral element mesh) and then to perform the waveform modeling (1) to obtain the wavefield (e). The second option is to use homogenization to compute the effective elastic model (c) from (a) valid up to f_{\max} . (c) is fully anisotropic but smooth (without being spatially uniform), which leads to a simple mesh (d). Finally the waveform modeling (2) is performed with the same numerical solver as the one used for waveform modeling (1) to obtain the wavefield (f), but at a lower numerical cost than (e). For a small enough ε_0 (see text), (f) converges to (e). The objective of this work is to present an alternate scheme to finite elements for the homogenization step allowing one to go from (a) to (c). The homogenization of this particular example (the Marmousi model) is detailed in [1].

for discontinuous media, they are not straightforward to implement (even if the finite element method is a very well-known process with many open libraries already there, it is still challenging for large problems) and requires a meshing procedure using software and skills that many seismologists do not have. For many applications, such as the case of continuous models (from the output of a full waveform inversion or a geological model given as pixelized, for instance), avoiding a finite element solver and the associated mesh generation is appealing. To solve the cell problem, as an alternative solution to finite elements, we propose to use the Fast Fourier Transform (FFT) iterative algorithm developed in [11] that we name Fast Fourier Homogenization (FFH). The boundary conditions of the cell problem can be periodic, even in the non-periodic case, making the FFT algorithm very well suited for such a problem. This algorithm has many advantages: it is simple to implement, it strongly relies on 2-D or 3-D FFT for which highly optimized packages exist and the associated regular mesh is trivial to generate. It is nevertheless, in principle, limited to continuous elastic media. To summarize, the objective of the paper is to propose an alternative approach to the finite element method to solve the cell problem of the non-periodic homogenization method allowing one to go from model (a) to (c) on Fig. 1.

The article is organized as follows: we first recall the main and useful non-periodic homogenization results. Then the FFT algorithm [11] and its application to the non periodic homogenization for elastic waves is presented, followed by the parallelization and the nested homogenization aspects. Finally, we show that the FFH gives very good results and, in many situations, it can even be used for discontinuous media with reasonable accuracy and

cost.

2. The homogenization algorithm to obtain the effective elastic tensor as well as source and receiver correctors

We give here a summary of the technique described in [1, 2, 3]. A reader interested in a more detailed and self-consistent justification of the method should refer to these articles.

In the following, we assume that the seismic source \mathbf{s} is such that it has a maximum frequency f_{\max} , making sure that, in the far-field, the wavefield has a minimum wavelength λ_{\min} . We consider an elastic medium described by its density and elastic tensor $(\rho(\mathbf{x}), \mathbf{c}(\mathbf{x}))$ at any point \mathbf{x} of the physical domain Ω . No assumption is made on the spatial variation of the density and of the elastic tensor; in particular, no periodicity, natural scale separation or any kind of spatial statistical invariance is assumed. This is an important point as realistic geological media are deterministic, heterogeneous at all scales and obviously non-periodic. The elasto-dynamic equations in Ω are:

$$\rho \partial_{tt} \mathbf{u} - \nabla \cdot \boldsymbol{\sigma} = \mathbf{s} \quad (3)$$

$$\boldsymbol{\sigma} = \mathbf{c} : \boldsymbol{\epsilon}(\mathbf{u}) \quad (4)$$

associated with the appropriate boundary conditions on $\partial\Omega$, where $\mathbf{u}(\mathbf{x}, t)$ is the ground displacement for $t \in [0, T]$, T the signal duration, $\boldsymbol{\sigma}(\mathbf{x}, t)$ the stress and $\boldsymbol{\epsilon}(\mathbf{u}) = \frac{1}{2}(\nabla \mathbf{u} + {}^t \nabla \mathbf{u})$ the strain. The homogenization technique, at order 0, aims to approximate the above original problem with the following effective equations

$$\rho^{\varepsilon_0} \partial_{tt} \mathbf{u}^{\varepsilon_0} - \nabla \cdot \boldsymbol{\sigma}^{\varepsilon_0} = \mathbf{s} \quad (5)$$

$$\boldsymbol{\sigma}^{\varepsilon_0} = \mathbf{c}^{\varepsilon_0} : \boldsymbol{\epsilon}(\mathbf{u}^{\varepsilon_0}) \quad (6)$$

still with the appropriate boundary conditions on $\partial\Omega$ (we describe issues associated with boundary conditions later on in this section), where $(\rho^{\varepsilon_0}(\mathbf{x}), \mathbf{c}^{\varepsilon_0}(\mathbf{x}))$ are the ε_0 effective density and elastic parameters, $\mathbf{u}^{\varepsilon_0}(\mathbf{x}, t)$ the order 0 homogenized displacement and $\boldsymbol{\sigma}^{\varepsilon_0}(\mathbf{x}, t)$ the average of the order 0 homogenized stress. Note that, unlike what is usually found in many homogenization processes, the effective properties are here not spatially uniform and still depend upon the space variable \mathbf{x} . The ε_0 parameter is a user defined parameter which controls the level of detail exhibited by the effective medium with respect to λ_{\min} :

$$\varepsilon_0 = \frac{\lambda_0}{\lambda_{\min}}. \quad (7)$$

λ_0 defines the value below which all scales are considered as small scales (also named “fast scales” or microscopic scales) and above which all scales are considered as large scales (also named “slow scales” or macroscopic scales). ε_0 is different from the classical small parameter ε used in periodic two scale homogenization [3]. We can interpret ε_0 in the following manner: an $\varepsilon_0 < 1$ means that the effective solutions will keep more details than λ_{\min} in the elastic model description, and $\varepsilon_0 > 1$ means that it will be smoother than λ_{\min} . In practice, for most geophysical applications, having ε_0 lying between 0.25 and 0.5 would be a good choice, but it all depends on the desired accuracy, on the signal duration and on the particular elastic model in consideration. The only point guaranteed by the method is the convergence rate with ε_0 (see below).

To the order 1, the relation between the true displacement \mathbf{u} and the homogenized displacement $\mathbf{u}^{\varepsilon_0}$ is

$$\mathbf{u}(\mathbf{x}, t) = \mathbf{u}^{\varepsilon_0}(\mathbf{x}, t) + \varepsilon_0 \boldsymbol{\chi}^{\varepsilon_0}(\mathbf{x}, \mathbf{x}/\varepsilon_0) : \boldsymbol{\epsilon}(\mathbf{u}^{\varepsilon_0})(\mathbf{x}, t) + O(\varepsilon_0), \quad (8)$$

where $\chi^{\varepsilon_0}(\mathbf{x}, \mathbf{y})$ is the first order corrector and accounts for the effect of structures local to the receiver on the recorded displacement (the site effect). χ^{ε_0} depends on two independent space variables: the regular space position \mathbf{x} and the microscopic scale variable \mathbf{y} . In practice, only $\mathbf{y} = \mathbf{x}/\varepsilon_0$ is useful. This two variable formulation is an important aspect of the two scale homogenization theory and one can refer to, for example, [4] for a deeper explanation. In practice, (8) is often valid up to $O(\varepsilon_0^2)$ [1]. Similarly, we have

$$\boldsymbol{\sigma}(\mathbf{x}, t) = \mathbf{H}^{\varepsilon_0}(\mathbf{x}, \mathbf{x}/\varepsilon_0) : \boldsymbol{\epsilon}(\mathbf{u}^{\varepsilon_0})(\mathbf{x}, t) + O(\varepsilon_0), \quad (9)$$

$$\boldsymbol{\epsilon}(\mathbf{x}, t) = \mathbf{G}^{\varepsilon_0}(\mathbf{x}, \mathbf{x}/\varepsilon_0) : \boldsymbol{\epsilon}(\mathbf{u}^{\varepsilon_0})(\mathbf{x}, t) + O(\varepsilon_0), \quad (10)$$

where $\mathbf{H}^{\varepsilon_0}(\mathbf{x}, \mathbf{y})$ and $\mathbf{G}^{\varepsilon_0}(\mathbf{x}, \mathbf{y})$ are the stress and strain concentrators and are also linked to site effects. Computing $(\rho^{\varepsilon_0}, \mathbf{c}^{\varepsilon_0})$ as well as χ^{ε_0} , $\mathbf{H}^{\varepsilon_0}$ and $\mathbf{G}^{\varepsilon_0}$ for non-periodic media is the main contribution of [1, 2, 3]'s work. Before moving forward in the description of the method, let us clarify and insist on some aspects of the non-periodic homogenization that might be confusing for readers familiar with other homogenization and upscaling techniques:

- No particular assumption is made on the spatial variations of the original media $(\rho(\mathbf{x}), \mathbf{c}(\mathbf{x}))$. In particular, no periodicity, no statistical invariance or any natural scale separation is assumed. As a consequence, most of the two scale periodic homogenization, stochastic homogenization or numerical homogenization techniques do not apply to our context. In the examples shown in Sec. 5, we use randomly generated elastic models, but each time, only one model realization is made. These elastic models are therefore used as deterministic models and shall not be interpreted as part of a stochastic process;

- an artificial scale separation is introduced by the user and its location in the wave-number domain is specified by the ε_0 value (see above). This allows one to define fast and slow scales, even if the medium doesn't present any natural scale separation;
- the effective media $(\rho^{\varepsilon_0}(\mathbf{x}), \mathbf{c}^{\varepsilon_0}(\mathbf{x}))$ is not uniform over the domain and depends upon \mathbf{x} . It is nevertheless small-scale free (it is “smoother” than the original medium, but usually still oscillates faster than λ_{\min}). The effective medium also depends upon the artificial scale separation through ε_0 . The smallest is ε_0 the more “detailed” is the effective medium with respect to λ_{\min} . As a consequence, the user can decide to keep more or less details from the original medium in the effective one in order to tune the accuracy of the wave propagation modeling;
- the non-periodic homogenization process is based on a cell problem (see below) formally similar to the cell problem present in all homogenization methods. Nevertheless, it is not related to any spatial cell as no periodic cell or any representative cell exists in our media. In the non-periodic homogenization case, the “cell” is the whole medium, or large chunks of the whole medium when a parallel implementation is needed (see Sec. 4.1). If periodic boundary conditions are still used for the cell problem, they are not related to any medium periodicity and have no impact of the final result;

To practically materialize the user defined ε_0 scale separation, we need to introduce a low-pass filter operator \mathcal{F}^{k_0} , such that, for any quantity $g(\mathbf{x})$, $\mathcal{F}^{k_0}(g)(\mathbf{x})$ does not contain any spatial variation faster than $\lambda_0 = 2\pi/k_0$.

This low-pass filter operator can be written as

$$\mathcal{F}^{k_0}(g)(\mathbf{x}) = (w^{k_0} * g)(\mathbf{x}), \quad (11)$$

where $*$ is the spatial convolution and w^{k_0} is the filter wavelet.

Going back to the non-periodic homogenization problem, the different steps allowing to build the upscaling operator \mathcal{H}^{k_0} such that

$$\mathbf{c}^{\varepsilon_0} = \mathcal{H}^{k_0}(\mathbf{c}) \quad (12)$$

and to find the correctors are the following:

- Step 1. We first solve the so-called cell problem to find the initial guess correctors $\chi_s^{lm}(\mathbf{x})$. It consists in solving the following elasto-static set of problems (3 in 2-D, 6 in 3-D) in Ω :

$$\begin{aligned} \nabla \cdot \mathbf{c} : \boldsymbol{\epsilon}(\chi_s^{lm}) &= \mathbf{F}^{lm}, \\ \mathbf{F}^{lm} &= -\nabla \cdot (\mathbf{c} : (\mathbf{e}_l \otimes \mathbf{e}_m)) \end{aligned} \quad (13)$$

with periodic boundary conditions on $\partial\Omega$, where the \mathbf{e}_i , $i \in \{1\dots d\}$, are the Cartesian unit basis vectors.

- Step 2. Once the initial corrector guess χ_s^{lm} is obtained, we compute

$$\begin{aligned} [\mathbf{G}_s]_{ijklm}(\mathbf{x}) &= \frac{1}{2} (\delta_{il}\delta_{jm} + \delta_{jl}\delta_{im}) + [\boldsymbol{\epsilon}(\chi_s^{lm})]_{ij}, \\ \mathbf{H}_s(\mathbf{x}) &= \mathbf{c} : \mathbf{G}_s. \end{aligned} \quad (14)$$

The ε_0 effective tensor can be directly obtained as

$$\mathbf{c}^{\varepsilon_0}(\mathbf{x}) = \mathcal{F}^{k_0}(\mathbf{H}_s) : \mathcal{F}^{k_0}(\mathbf{G}_s)^{-1}(\mathbf{x}). \quad (15)$$

At this stage the upscaling operator \mathcal{H}^{k_0} , defined in Eq. 12, is known.

The effective density is simply

$$\rho^{\varepsilon_0}(\mathbf{x}) = \mathcal{F}^{k_0}(\rho)(\mathbf{x}); \quad (16)$$

- Step 3. Finally, the stress and strain concentrators are obtained as

$$\begin{aligned}\mathbf{G}^{\varepsilon_0}(\mathbf{x}, \mathbf{y}) &= (\mathcal{F}^{k_0}(\mathbf{G}_s^{\varepsilon_0})(\mathbf{x}) + (\mathbf{G}_s^{\varepsilon_0} - \mathcal{F}^{k_0}(\mathbf{G}_s^{\varepsilon_0}))(\varepsilon_0 \mathbf{y})) : \mathcal{F}^{k_0}(\mathbf{G}_s^{\varepsilon_0})^{-1}(\mathbf{x}), \\ \mathbf{H}^{\varepsilon_0}(\mathbf{x}, \mathbf{y}) &= (\mathcal{F}^{k_0}(\mathbf{H}_s^{\varepsilon_0})(\mathbf{x}) + (\mathbf{H}_s^{\varepsilon_0} - \mathcal{F}^{k_0}(\mathbf{H}_s^{\varepsilon_0}))(\varepsilon_0 \mathbf{y})) : \mathcal{F}^{k_0}(\mathbf{G}_s^{\varepsilon_0})^{-1}(\mathbf{x})\end{aligned}\quad (17)$$

and the first order corrector $\boldsymbol{\chi}^{\varepsilon_0}(\mathbf{x}, \mathbf{y})$ is obtained solving, for each \mathbf{x} (fixed),

$$\boldsymbol{\nabla} \boldsymbol{\chi}^{\varepsilon_0}(\mathbf{x}, \mathbf{y}) + {}^T \boldsymbol{\nabla} \boldsymbol{\chi}^{\varepsilon_0}(\mathbf{x}, \mathbf{y}) = 2(\mathbf{G}^{\varepsilon_0}(\mathbf{x}, \mathbf{y}) - \mathbf{I}), \quad (18)$$

where the $\boldsymbol{\nabla}$ operators here applies on the \mathbf{y} variable and \mathbf{I} is the identity operator.

At this stage, everything is ready: the effective medium is known and can be used in our favorite wave equation solver. The local effect operators are also ready to be applied as a post-process to the output wavefield computed by wave equation solver. The effectiveness of the non-periodic homogenization summarized here has been demonstrated in, for example, [1] and we do not propose any new validation of the method here. Only an alternative cell problem solver is proposed and tested in the following.

In step 1, the cell problem (13) is solved on the whole domain $\boldsymbol{\Omega}$ at once with periodic boundary conditions. This is very different to what is often done with numerical homogenization techniques (e.g. the Heterogeneous Multiscale Method (HMM) [12]) which are based on many small cell problems sampling the domain and where the cell boundary conditions are often a difficulty. In our case, the boundary conditions have an effect on the final result only in the vicinity of the domain boundary, in an area whose extent scales as λ_0 [1]. To avoid any spurious effect near $\partial\boldsymbol{\Omega}$, a buffer with constant

elastic properties can be added around Ω . In practice, Ω can be large, leading to a cell problem that does not fit in the memory of a single computer. In such a case, splitting the domain into multiple smaller sub-domains can be necessary. To prevent error at sub-domain junctions, the sub-domains are then overlapping each other over an area of extent proportional to λ_0 . This aspect is discussed in more details in Sec. 4.1.

Finally let us mention that nothing has been said about the boundary conditions on the original wave equation problem (like the Neumann condition at the free surface), which should not be confused with the boundary conditions of the cell problem, and how they can be modified by the homogenization process. They usually deserve a special treatment as described in [7] and [13].

3. The iterative algorithm to solve the cell problem

The most computationally challenging step of the previous work-flow is solving the cell problem (13). The only case where it can be solved analytically is the layered case [14, 3], which allows one to obtain the well-known Backus’s result [5]. In general, we rely on a finite element solver to solve (13), as is done in [1] and [2]. As mentioned in the introduction, a finite element solver and its associated mesh are difficult to set up and consequently makes the possibility to avoid them appealing. For many applications, such a method adapted to discontinuous models is not necessary. Indeed, many elastic models are given as continuous (for example: outputs of full waveform inversions, or pixelized geological models). To solve the cell problem (13) in the continuous model case, we propose using the FFT iterative algorithm [11], initially introduced in the structural mechanics context. This method is

based on a simple regular mesh and is therefore very easy to use. Moreover, it is perfectly suited for large domains with periodic boundary conditions. We summarize here their results and one should refer to the original article for more details.

3.1. The periodic Lippman-Schwinger equation

We first start with an auxiliary problem for which a solution can be easily computed:

$$\begin{aligned}\nabla \cdot \mathbf{c}^0 : (\boldsymbol{\epsilon}(\mathbf{v})) &= \mathbf{F} , \\ \mathbf{F} &= -\nabla \cdot \boldsymbol{\tau}\end{aligned}\tag{19}$$

in Ω , \mathbf{v} being Ω -periodic and \mathbf{c}^0 constant over Ω .

For such a problem, we can write

$$\boldsymbol{\epsilon}(\mathbf{v})(\mathbf{x}) = \int_{\Omega} \Gamma^0(\mathbf{x} - \mathbf{x}') : \boldsymbol{\tau}(\mathbf{x}') d\mathbf{x}' .\tag{20}$$

where Γ^0 is the periodic Green function associated with \mathbf{c}^0 . For a simple constant \mathbf{c}^0 tensor, Γ^0 can be computed analytically as shown by [11].

The above auxiliary problem (19) and its solution (20) can be used to solve the following more general problem:

$$\begin{aligned}\nabla \cdot \mathbf{c} : (\boldsymbol{\epsilon}(\mathbf{v})) &= \mathbf{F} , \\ \mathbf{F} &= -\nabla \cdot (\mathbf{c} : \mathbf{E}) ,\end{aligned}\tag{21}$$

where the tensor $\mathbf{c}(\mathbf{x})$ now depends upon \mathbf{x} and \mathbf{E} is a prescribe constant strain over Ω . The last equations can be re-written as

$$\begin{aligned}\nabla \cdot \mathbf{c}^0 : (\boldsymbol{\epsilon}(\mathbf{v})) &= \mathbf{F}' , \\ \mathbf{F}' &= -\nabla \cdot ((\mathbf{c} - \mathbf{c}^0) : \boldsymbol{\epsilon}(\mathbf{v}) + \mathbf{c} : \mathbf{E}) .\end{aligned}\tag{22}$$

Using the auxiliary problem we find the Lippman-Schwinger equation for $\boldsymbol{\epsilon}(\mathbf{v})$:

$$\boldsymbol{\epsilon}(\mathbf{v}) = \mathbf{E} + \int_{\Omega} \boldsymbol{\Gamma}^0(\mathbf{x} - \mathbf{x}') : ((\mathbf{c} - \mathbf{c}^0) : \boldsymbol{\epsilon}(\mathbf{v})) (\mathbf{x}') d\mathbf{x}'. \quad (23)$$

To solve the above implicit equation, [11] rely on the iterative scheme described in the next subsection.

3.2. The algorithm

To solve efficiently (23), [11] propose an iterative scheme based on the analytical expression of $\boldsymbol{\Gamma}^0$ for a constant isotropic medium \mathbf{c}^0 (defined by its Lamé coefficients λ^0 and μ^0). We first define, for any $g(\mathbf{x})$, its Fourier transform $\hat{g}(\mathbf{k})$:

$$\hat{g}(\mathbf{k}) = \mathcal{F}(g)(\mathbf{k}) = \int_{\mathbb{R}^d} g(\mathbf{x}) e^{i\mathbf{k}\cdot\mathbf{x}} d\mathbf{x}, \quad (24)$$

where \mathbf{k} is the wave-number vector. Using the inverse transformation, we can reconstruct g from \hat{g} :

$$g(\mathbf{x}) = \mathcal{F}^{-1}(\hat{g})(\mathbf{x}) = \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} \hat{g}(\mathbf{k}) e^{-i\mathbf{k}\cdot\mathbf{x}} d\mathbf{k}. \quad (25)$$

The convolution theorem, which states that for any g and h

$$g * h(\mathbf{x}) = \int_{\mathbb{R}^d} g(\mathbf{x} - \mathbf{x}') h(\mathbf{x}') d\mathbf{x}' = \mathcal{F}^{-1}(\hat{g}(\mathbf{k}) \hat{h}(\mathbf{k})), \quad (26)$$

will be useful in the following. [11] show that the Green function of the auxiliary problem can be easily evaluated in the Fourier domain:

$$\hat{\Gamma}_{ijlm}^0(\mathbf{k}) = \frac{1}{4\mu^0|\mathbf{k}|^2} (\delta_{li}k_mk_j + \delta_{mi}k_lk_j + \delta_{lj}k_mk_i + \delta_{mj}k_lk_i) - \frac{\mu^0 + \lambda^0}{\mu^0(\lambda^0 + 2\mu^0)} \frac{k_ik_jk_lk_m}{|\mathbf{k}|^4}. \quad (27)$$

Applying the convolution theorem (26), [11] propose the following iterative scheme to solve (22):

- Initialization

$$\boldsymbol{\epsilon}^0(\mathbf{x}) = \mathbf{E} \quad \forall \mathbf{x} \in \Omega \quad (28)$$

$$\boldsymbol{\sigma}^0(\mathbf{x}) = (\mathbf{c} : \boldsymbol{\epsilon}^0)(\mathbf{x}) \quad \forall \mathbf{x} \in \Omega \quad (29)$$

- from iteration i to $i + 1$

1. $\hat{\boldsymbol{\sigma}}^i = \mathcal{F}(\boldsymbol{\sigma}^i)$

- 2.

$$\begin{cases} \hat{\boldsymbol{\epsilon}}^{i+1}(\mathbf{k}) &= \hat{\boldsymbol{\epsilon}}^i(\mathbf{k}) - (\hat{\Gamma}^0 : \hat{\boldsymbol{\sigma}}^i)(\mathbf{k}) \quad \forall \mathbf{k} \neq \mathbf{0} \\ \hat{\boldsymbol{\epsilon}}^{i+1}(\mathbf{0}) &= \mathbf{E} \end{cases}$$

3. $\boldsymbol{\epsilon}^{i+1} = \mathcal{F}^{-1}(\hat{\boldsymbol{\epsilon}}^{i+1})$

4. $\boldsymbol{\sigma}^{i+1}(\mathbf{x}) = (\mathbf{c} : \boldsymbol{\epsilon}^{i+1})(\mathbf{x}) \quad \forall \mathbf{x} \in \Omega$

5. convergence test

We stop the iterations when the convergence test is met, that is $|\int_{\Omega} \boldsymbol{\sigma}^{i+1} - \boldsymbol{\sigma}^i d\mathbf{x}|$ has reached a small enough value.

The actual value of μ^0 and λ^0 is not crucial, but following [11] we observed that

$$\lambda^0 = \frac{1}{2} \left(\inf_{\mathbf{x} \in \Omega} \lambda + \sup_{\mathbf{x} \in \Omega} \lambda \right) \quad (30)$$

$$\mu^0 = \frac{1}{2} \left(\inf_{\mathbf{x} \in \Omega} \mu + \sup_{\mathbf{x} \in \Omega} \mu \right) \quad (31)$$

give the best convergence rate. In the anisotropic case, in (30) and (31), we choose $\lambda(\mathbf{x})$ and $\mu(\mathbf{x})$ from the closest isotropic elastic tensor $\mathbf{c}_{\text{iso}}(\mathbf{x})$ to $\mathbf{c}(\mathbf{x})$ for each \mathbf{x} obtained following the projection method of [15]. Let us note that it is remarkable that such simple homogeneous media Green functions

can be employed to solve efficiently the cell problem in very complex media, with very large elastic heterogeneity contrasts (even with orders of magnitude contrasts as shown by [11]).

3.3. Application to the non-periodic homogenization

The above algorithm can be directly applied to the step 1 of the non-periodic homogenization with

$$\mathbf{E} = \frac{1}{2} (\mathbf{e}_l \otimes \mathbf{e}_m + \mathbf{e}_m \otimes \mathbf{e}_l) \quad (32)$$

for each $(m, l) \in \{1, \dots, d\}^2$ (3 possibilities in 2-D and 6 in 3-D). Doing so, the generic problem (21) is indeed the same as the cell problem (13). For a given domain Ω and elastic tensor \mathbf{c} , finding $\mathbf{c}^{\varepsilon_0}$ therefore implies using 3 times in 2-D and 6 times in 3-D the iterative algorithm presented in Sec. 3.2 to obtain the initial guess corrector χ_s^{lm} . Once this is done, finding the effective medium using (15), the final concentrators $\mathbf{H}^{\varepsilon_0}$ and $\mathbf{G}^{\varepsilon_0}$ using (17) and final corrector χ^{ε_0} using (18) is trivial. Moreover the spatial low-pass filtering (11) can be efficiently computed using the convolution theorem, such as, for any $g(\mathbf{x})$

$$\mathcal{F}^{k_0}(g)(\mathbf{x}) = \mathcal{F}^{-1}(\hat{w}^{k_0} \hat{g}) . \quad (33)$$

The application of the FFT algorithm to solve the cell problem of the non-periodic homogenization for the wave equation is the FFH method.

4. Numerical considerations

Implementing the iterative algorithm presented in Sec. 3.2 is relatively simple and can be very effective when based on a standard FFT package, like the FFTW library [16] we are using here. Compared to the original

application of the iterative algorithm [11], the non-periodic homogenization is more memory intensive. This is because the $\mathbf{G}_s^{\varepsilon_0}$ tensor, a 9 dimension tensor in 2-D and a 21 dimension tensor in 3-D, has to be stored at some point in the computation. This intensive memory usage makes the parallel implementation quickly necessary for realistic cases, especially in 3-D.

4.1. Parallel implementation

When the domain Ω is large enough such that the cell problem does not fit in a single shared memory processor, a parallel implementation becomes necessary. To solve the cell problem in parallel, we could design a solution based on a domain decomposition, but still on a global computation implying the whole domain Ω at once. This would be technically possible (knowing that most FFT packages allow such a parallel implementation), but it would involve a large volume of communication between processors and it would necessarily impose a large parallel computer for a large domain (the whole problem has to fit in the distributed memory at once).

Instead, we propose an embarrassingly parallel implementation, in which sub-domains are considered as independent from each other, allowing to treat them sequentially if necessary, which is an important advantage compared to the classical domain decomposition. Indeed, with such a scheme, large domains can be homogenized even with a small computer by solving the cell problem on each sub-domain one by one. It would certainly take some time compared to a large computer, but it would be possible. To build such a scheme, we first decompose Ω in N_d sub-domains Ω_i , such that $\Omega = \cup_{i=1}^{N_d} \Omega_i$, but, unlike many classical domain decompositions, the sub-domains Ω_i overlap with their neighbors. The overlapping areas are named

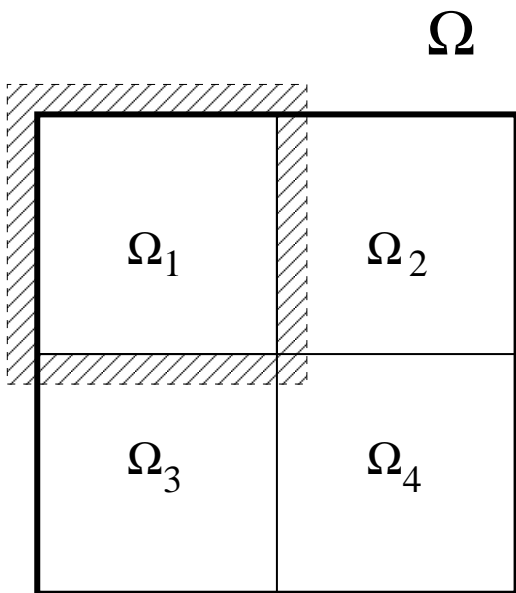


Figure 2: Sketch showing a Ω domain split into 4 sub-domains. The buffer for sub-domain Ω_1 is shown in hatch. A sub-domain is built of the sub-part of Ω and the buffer around it.

“buffers”. Note that the size of the sub-domains is determined by the shared memory size of the processors (the larger is the memory, the larger sub-domains are) and is only indirectly linked to the scale separation length λ_0 through the buffers. An example of such a decomposition is given in Fig. 2, where a domain Ω is split into 4 sub-domains. For the sub-domain Ω_1 , the buffer is shown and is overlapping with the other 3 sub-domains. Note that in that case, Ω_1 is touching $\partial\Omega$, which imposes to define values for the elastic properties outside Ω . This of course influences the result of the homogenization near $\partial\Omega$ and, as mentioned earlier, nothing is said here about the boundary conditions on the original problem (like the Neumann condition at the free surface). They usually deserve a special treatment as described in [7] and [13]. The important point is the width of the buffer. The

boundary effects decrease exponentially away from the sub-domain boundary [17, 18]. Nevertheless, they are spread by the low-pass filter applied in (15) and (17). To avoid spurious effects of the domain decomposition, the buffer width must be larger and proportional to half of the support of the filter wavelet w^{k_0} (see eq. 11). If the buffer is wide enough, no spurious effect of the domain decomposition appears as it will be shown in the validation test section.

4.2. Nested homogenization

In practice, for large enough domains, the buffer can quickly become a problem, especially in 3-D. Indeed, the memory required to solve the cell problem in the buffer part can be large compared to the one involved for the sub-domain itself, which severely degrades the efficiency of the parallel algorithm. To understand the problem, we first assume that we need to homogenize, with $\varepsilon_0 = 2\pi/(k_0\lambda_{\min})$, a given elastic model whose Fourier spectrum has a maximum wave-number k_{\max} . In order to sample correctly the medium, the distance between two FFT sampling points is proportional to $1/k_{\max}$ (the shortest sampling distance corresponds to the Nyquist wave-number, π/k_{\max}) and therefore the memory necessary to solve the cell problem scales as $(k_{\max})^d$. Furthermore, we note that the support of w^{k_0} , and therefore the buffer size, is directly proportional to $1/k_0$, which implies that, the smaller is k_0 , the wider is the buffer. In the case of $k_0 \ll k_{\max}$, it may happen that discretizing the buffer area only gives an homogenization problem whose size does not fit into a processor memory. In such a case, dividing the domain Ω in even smaller sub-domains does not help anymore and consequently another solution has to be found. To circumvent this difficulty, we




Figure 3: Sketch of a domain Ω homogenized in 3 nested steps, with two intermediate values of ε_0 : $\varepsilon_{0,2} < \varepsilon_{0,1} < \varepsilon_0$. First, to homogenize in $\varepsilon_{0,2}$, the domain is divided into 16 sub-domains (left). Then, to homogenize in $\varepsilon_{0,1}$, the domain is divided into 4 sub-domains (middle). Finally the homogenization is carried out to the ε_0 level on the whole domain at once. It can be noticed that the buffer (in hatch) is wider with increasing $\varepsilon_{0,i}$.

propose a method for which the main idea is to perform the homogenization in several nested steps. We first homogenize with a large $k_{0,N}$ (potentially much larger than the target k_0), therefore using a small buffer which makes the domain decomposition efficient. The resulting effective tensor \mathbf{c}^N has a Fourier spectrum maximum wave-number proportional to $k_{0,N}$, smaller than k_{\max} . Homogenizing \mathbf{c}^N can now be done with a new wave-number $k_{0,N-1}$, larger than $k_{0,N}$. This second homogenization implies a larger buffer (because $k_{0,N-1}$ is smaller than $k_{0,N}$) but also a sampling sparser than for the previous step (because $k_{0,N}$ is smaller than k_{\max}), making the computation possible. The operation is then repeated until the target k_0 is reached. Performing the homogenization in several nested steps with a series of ε_0 ($\varepsilon_0 > \varepsilon_{0,1} > \dots > \varepsilon_{0,i} > \dots > \varepsilon_{0,N}$) allows to keep the buffer small at each step. A sketch illustrating this idea is given in Fig. 3 and a formal description of the nested algorithm and its consequences on correctors is given in Appendix A.

5. Validation tests

All the presented tests are performed in a 2-D geometry. Tests for the 3-D homogenization will be presented in a separate paper. The objective here is to validate the FFH method and not the non-periodic homogenization itself, which has already been tested in [1, 2, 3]. The convergence of the homogenization method will nevertheless be used as an indirect test to check the FFH algorithm. Note that two very different convergences will be discussed: the convergence of the iterative cell problem with the number of iterations (see Sec. 3) and the convergence of the non-periodic homogenization method with ε_0 (see Eq. (8)). They should not be confused.

5.1. Tests in continuous media

In this section, the method is tested in two 2-D heterogeneous media presented in Fig. 4. The media are made of a $40 \times 40 \text{ km}^2$ domain of constant density and elastic properties, with an inner $30 \times 30 \text{ km}^2$ domain of heterogeneous density and elastic properties. For the inner square, two media are designed. Each of them correspond to a test:

- Test (a) has a k^{-2} Fourier amplitude spectrum plus a high wave-number texture for the S-wave velocity V_s (see Fig. 4.a and 4.c).
- Test (b) has a $k^{-1/2}$ Fourier amplitude spectrum for V_s (see Fig. 4.b and 4.d).

The two models are generated randomly, but they are used in a deterministic way: a single realization of each model is performed and they are used as deterministic models afterward. This process is just a simple way to generate

challenging models for the method. The first model spectrum is chosen to obtain a slow variation and a fast variation above it in order to check potential spurious effects of the buffer when dealing with slow variations. The second model is designed to be very heterogeneous at all scales, which is useful to check that the convergence of the homogenization method as a function of ε_0 is not broken.

For the two models, V_p and ρ are scaled from the heterogeneous V_s . In the outer square, we have $V_p = 5 \text{ km s}^{-1}$, $V_s = 3 \text{ km s}^{-1}$ and $\rho = 3000 \text{ kg m}^{-3}$. For the experiments, an explosive source is triggered in $\mathbf{x}_s = (1 \text{ km}, 20 \text{ km})$, with a Ricker (second derivative of a Gaussian) source time function of central frequency of 1.1 Hz and maximum frequency of about 3 Hz , such that the shortest wavelength λ_{\min} outside the inner square is about 1 km . A reference solution is computed with the spectral element method (SEM, see, e.g. [19, 20]) with a mesh dense enough to capture all the complexity of the medium (an element size of $100 \times 100 \text{ m}^2$ and a polynomial degree equal to 7 in each direction have been chosen). Absorbing Perfectly Matched Layer (PML as described in [21]) are used around Ω . Energy snapshots of the reference computations are plotted in Fig. 4.e and 4.f.

We first test the embarrassingly parallel and nested algorithm of the FFH scheme. To do so, we first use the FFH on the whole domain at once (no parallelization). We use a 10^{-4} precision on the convergence test of the iterative algorithm. For the two models, the iterative scheme converges with less than 20 iterations. Many details on the convergence of the FFT iterative scheme are given in [11, 22]. Compared to the tests presented in the original works,

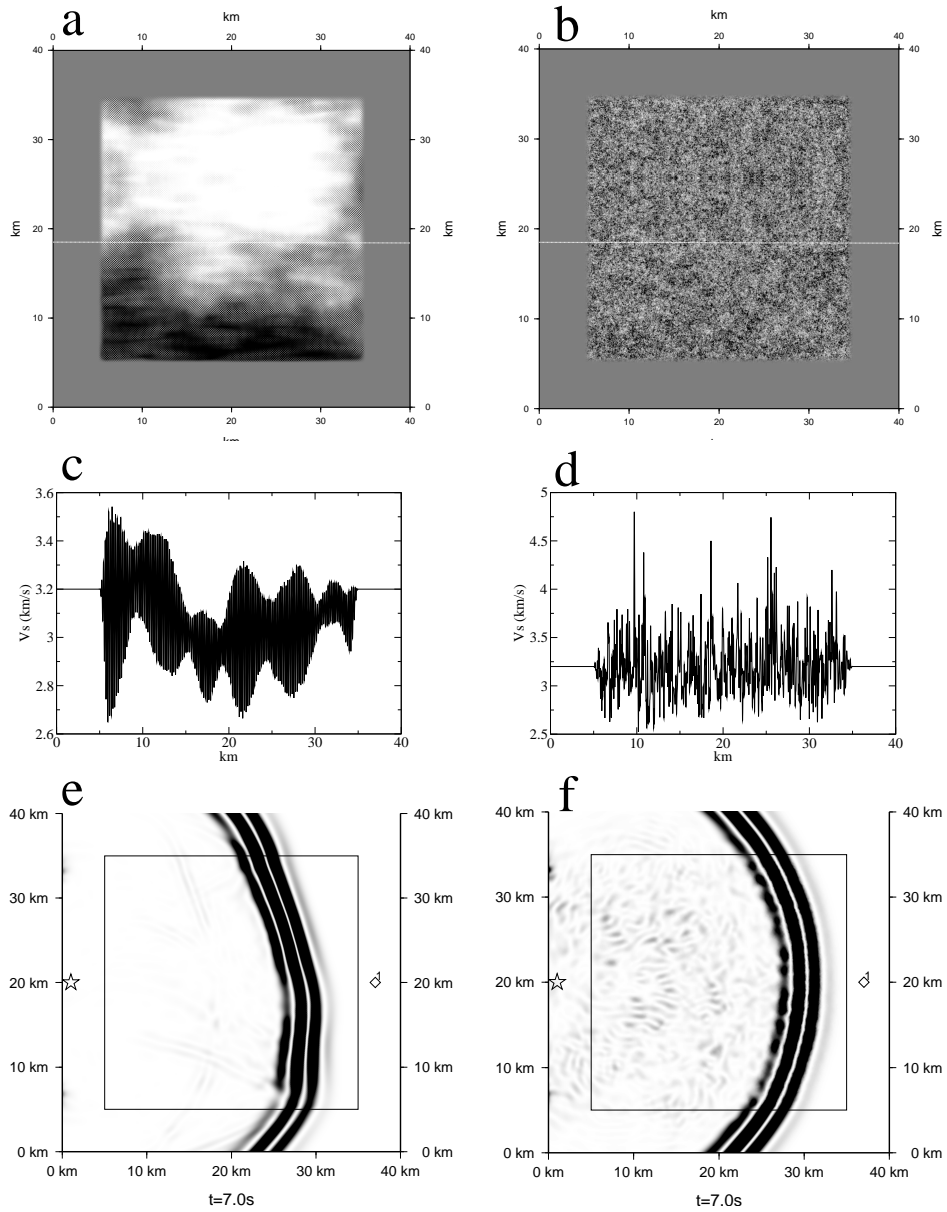


Figure 4: Continuous media used for the validation tests. In **a** and **b** is plotted the S-wave velocity V_s in gray scale (black is fast, white is slow) for test (a) and test (b) models, respectively. In **c** and **d** is plotted an horizontal cross-section at $y = 1850 \text{ m}$ of V_s from test (a) and (b) models, respectively (the position of the $y = 1850 \text{ m}$ cross-section horizontal line is plotted as a white line in **a** and **b**). In **e** and **f** are plotted energy snapshots for an explosion triggered in $\mathbf{x}_s = (1 \text{ km}, 20 \text{ km})$ (star) for $t = 7 \text{ s}$ in test (a) and (b) models, respectively.




Figure 5: Top graph: V_s cross-section at $y = 1850$ m in the original model (gray line) and in the $\varepsilon_0 = 0.5$ homogenized model (black line). Bottom graph: zoom around $x = 20$ km in the V_s cross-section at $y = 1850$ m in the original model (thin gray line), for the one step $\varepsilon_0 = 0.5$ homogenized model (black line), for the 4×4 sub-domain embarrassingly parallel algorithm with a (voluntarily) too small buffer (thin black dashed line) and the 4×4 sub-domain embarrassingly parallel algorithm with a correct buffer width (bold gray line). The nested algorithm has also been used (not shown as indistinguishable from the black line).

the elastic heterogeneity contrasts we are using here are not challenging (we have a maximum factor 4 for the ratio of the maximum versus minimum λ in the domain while [11, 22] tested up to several order of magnitude) and the algorithm is quickly converging without any difficulty. To challenge the parallel and nested schemes, we just check that the homogenization carried out on the whole domain Ω at once and a homogenization carried out in multiple steps (i.e. using several sub-domains) give the same results. Fig. 5 shows cross-sections of the S-wave velocity in the original medium and in the homogenized media for $\varepsilon_0 = 0.5$ (and $\lambda_{\min} = 1 \text{ km}$) computed in the whole domain at once, computed in 4×4 sub-domains with a too small buffer width, and finally computed in 4×4 sub-domains with a correct buffer width (a correct buffer width is, at least, half the size of the wavelet w^{k_0} support). If the buffer is too small, it can be seen that an error is introduced at the sub-domains junction (here around $x = 20 \text{ km}$). If the buffer width is correct, then the two methods give the same result. Finally, the nested method has been used in two steps (with $\varepsilon_{0,1} = 0.3$), starting from a 4×4 sub-domains decomposition. The result is the same as the one obtained in one step (no line has been plotted in Fig. 5 because the results are within the 10^{-4} precision and could not be distinguished from the one step solution). If this test shows that the different parallel and serial versions give the same result, it is not yet an accuracy test of the FFH scheme.

To go further in testing the FFH, because no reference solution exists for the homogenized model, we check that the convergence of the homogenized solution $\mathbf{u}^{\varepsilon_0}$ with ε_0 toward the reference solution is still the one observed in [1]. This is an indirect test, but if it is conclusive, we then will be able to

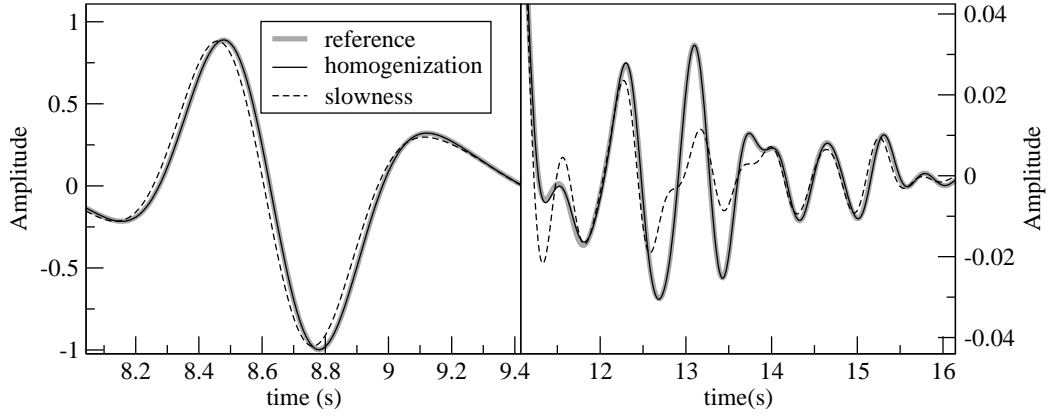


Figure 6: Horizontal component of the velocity for a receiver at $\mathbf{x} = (38\text{km}, 20\text{km})$ for the reference (gray line), for the $\varepsilon_0 = 0.5$ homogenized ($\dot{u}_1^{\varepsilon_0}$, black line) and for the “slowness” (dashed line) solutions for the test (a) elastic model. Two different time windows are considered: an early one (left plot) and a late one (right plot).

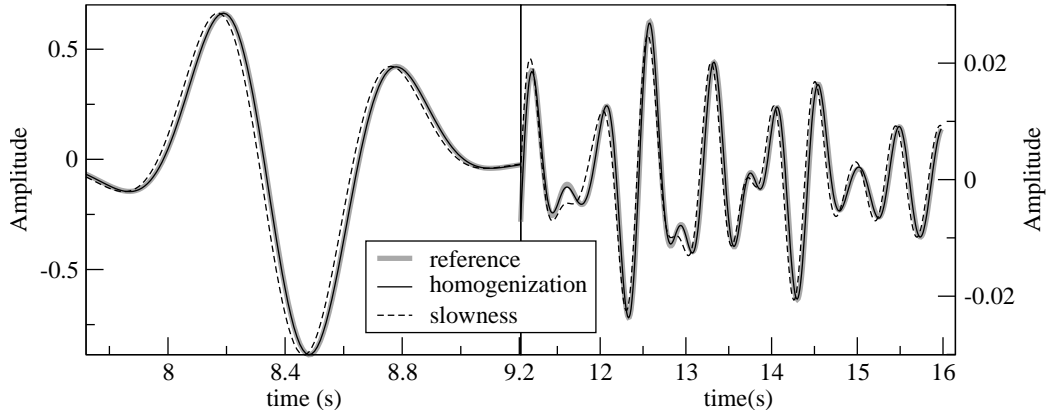


Figure 7: Horizontal component of the velocity for a receiver at $\mathbf{x} = (38\text{ km}, 20\text{ km})$ for the reference (gray line), for the $\varepsilon_0 = 0.5$ homogenized ($\dot{u}_1^{\varepsilon_0}$, black line) and for the “slowness” (dashed line) solutions for the test (b) elastic model. Two different time windows are considered: an early one (left plot) and a late one (right plot).

assess that the FFH method is accurate enough. To perform this test, the waveforms are computed with the SEM. Fig. 6, respectively Fig. 7, show the horizontal component for a receiver at $\mathbf{x} = (38 \text{ km}, 20 \text{ km})$ for the reference solution, for the $\varepsilon_0 = 0.5$ homogenized solution and the for “slowness” solution, computed for the test (a), respectively test (b), elastic models. The “slowness” solution is computed in a smooth model such that $\rho^{\varepsilon_0} = \mathcal{F}^{k_0}(\rho)$, $1/V_p^{\varepsilon_0} = \mathcal{F}^{k_0}(1/V_p)$ and $1/V_s^{\varepsilon_0} = \mathcal{F}^{k_0}(1/V_s)$ with the same k_0 as for the homogenization. The reason to show the “slowness” solution is to make sure that a simple naive effective model obtained by filtering the slowness does not perform as well as the homogenized solution, which makes our test relevant. It can be seen that the homogenization gives an accurate result for the ballistic wave as well as for the coda, which is not the case for the “slowness” solution. Finally, we look more closely at the convergence with ε_0 of the method in the test (b) model (we choose test (b) for its large amount of heterogeneities at all scales). To do so we define the error $E(\dot{\mathbf{u}})$ of a velocity field $\dot{\mathbf{u}}$ with respect to the reference solution $\dot{\mathbf{u}}^{\text{ref}}$ for a set of 20 receivers located in the inner square as:

$$E(\dot{\mathbf{u}}) = \sum_{i=1,20} \frac{\sqrt{\int_0^{t_{\max}} (\dot{\mathbf{u}} - \dot{\mathbf{u}}^{\text{ref}})^2(\mathbf{x}_i, t) dt}}{\sqrt{\int_0^{t_{\max}} (\dot{\mathbf{u}}^{\text{ref}})^2(\mathbf{x}_i, t) dt}} \quad (34)$$

with $t_{\max} = 24 \text{ s}$. Fig. 8 shows the above error for the homogenized and the “slowness” solutions as a function of ε_0 . It can be seen that the homogenization shows roughly a convergence in ε_0^2 in the range $[0.25, 1]$ whereas the “slowness” solution shows only a very slow convergence. Note that the reason why the homogenization does not show a better convergence for $\varepsilon_0 < 0.25$ is due to the fact that, below $\varepsilon_0 = 0.25$, the order 2 homogenization error is low




Figure 8: Error $E(\dot{\mathbf{u}}^{\epsilon_0})$ (see Eq. 34) for seismograms computed in the effective media obtained with the homogenization FFH tool (black square symbols) and with the “slowness” solutions (black circle symbols) as a function of ϵ_0 . The $y = \epsilon_0^2$ line has been plotted for comparison.

enough so that other errors, including the first order corrector, but also error from the SEM discretization and the fact that we have stopped the iterative scheme with a 10^{-4} precision, become significant.

In conclusion, these tests show very similar behavior for the accuracy and convergence rate to the one observed by [1] with a finite element algorithm, which is what was targeted. For this continuous model case, the algorithm shows good performances: for the two models, the homogenization could be run in about 90 seconds (among which 70 seconds were used to read and interpolate the reference model) using a 12 core CPU, to be compared with 20 minutes on 120 cores needed to compute the reference solution with the SEM.

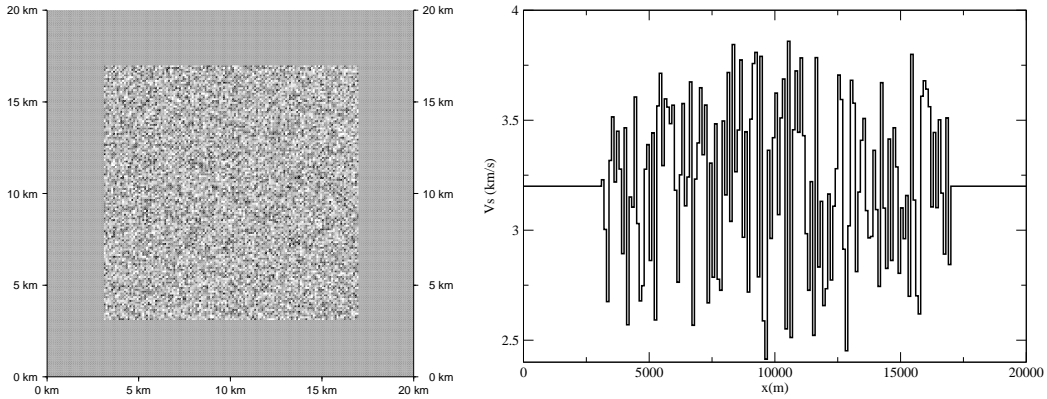


Figure 9: Left: S wave velocity V_s in Ω for the discontinuous test model (white colors correspond to low values and black colors to high values). Right: V_s horizontal cross-section at $y = 9550$ m as a function of x .

5.2. Tests in a discontinuous medium

In many cases, geological elastic models are discontinuous and it would be interesting to use the FFH method in such cases. Obviously, because Fourier series do not converge strongly at discontinuities (Gibbs phenomenon), we expect some difficulties in such cases. In this section, we do not address discontinuous media in general, but we treat one case and we show that it can work if some care is given.

We choose a similar random case to the one used by [1]. The domain Ω is a 20×20 km² domain of constant elastic properties (the same ones as in the previous section) with a 14×14 km² inner domain made of 140×140 little squares of constant elastic properties and density. The elastic properties and density of the little squares are generated randomly and independently of each others within $\pm 50\%$ of the outer square elastic values (see Fig. 9). Only one realization of the randomly generated model is used making it a deterministic model.

If, in the continuous model case, the FFT sampling can be easily estimated based on the spectral content of the elastic and density properties to be homogenized, in the discontinuous case, estimating the sampling rate is difficult. In our different tests, we found that the sampling rate depends upon the desired accuracy, the ε_0 value and the elastic and density contrasts at discontinuities. We first compute the homogenized model with 3 different samplings ($2.5 m$, $5 m$ and $10 m$) and compare the results with the one obtained with the finite element program used by [1]. Note that the finite element solution is not a reference solution here, but rather a comparison point. Indeed the finite element solution is also an approximate solution, but it is nevertheless obtained from a method adapted to discontinuous media. Interestingly, the sampling rate significantly affects the result of the effective model itself, as it can be seen on the cross-section presented in Fig. 10. Nevertheless, if different sampling rates have a significant effect on effective properties, their effects on seismograms computed in the effective media is much smaller. To quantify this observation, we use once again the SEM to compute a reference solution in the original model and solutions in each homogenized models. The source is the same explosion as for the previous tests (i.e. located in $\mathbf{x}_s = (1 km, 10 km)$). The signal duration is now $50 s$ long to let the error grow with time. No absorbing boundary has been added around the domain so that all the energy remains in Ω . We once again compute the error (34). As it can be seen in Tab. 1, the error computed for $\varepsilon_0 = 0.25$ using the finite element solver on the one hand, and the FFH method with the three smallest sampling rates on the other hand, are very similar. This confirms that, despite the visible effect of the sampling on the effective tensor, the




Figure 10: Top graph: horizontal cross-section at $y = 9550$ m in the S wave velocity of the effective model as a function of x , for $\varepsilon_0 = 0.5$, computed with the finite element scheme (FE) and the FFH scheme with three different sampling rates ($dx = 10$ m, $dx = 5$ m and $dx = 2.5$ m). Bottom graph: zoom in a part of the top graph.

	FE	$dx = 40m$	$dx = 20m$	$dx = 10m$	$dx = 5m$	$dx = 2.5m$
Error	0.048	0.171	0.089	0.067	0.042	0.038

Table 1: Error $E(\mathbf{u}^{\varepsilon_0})$, as defined by Eq. 34, for 20 receivers in the inner square and 50 s of signal, comparing the reference solution and the homogenized solution for $\varepsilon_0 = 0.25$, where the effective media has been computed with the finite element solver (FE) and the FFH solver using different sampling rates dx .

final effect on the homogenized seismograms is weak. This implies that the sampling rate of the FFH algorithm (if dense enough) mainly affects the high wave-number part (in the homogenization sense) of the effective properties, which have little effects on the homogenized seismograms.

The computation time of both FE and FFH strongly depends upon the sampling, number of elements and degree used. If two elements per little squares and a 1-D degree 6 are chosen for the FE method, then a $dx = 7m$ sampling has to be chosen for the FFH method to achieve the same final error on synthetic seismograms. In that case the FFH method is much faster than the finite element method: using a 12 core CPU, it took 40 s for the FFH method and 550 s for the finite element method to complete the computation.

Finally, we look at the corrector computed by the FFH algorithm in this discontinuous medium. To do so, in Fig. 11, we compare a sample of a cross-section in a component of the strain corrector ($G_{1111}^{\varepsilon_0}$) computed for $\varepsilon_0 = 0.25$ with the finite element algorithm and the FFH algorithm. For the finite element method, the interpolation is naturally performed based on the polynomial interpolation per element embedded in the method. For the FFH algorithm, the Fourier expansion should be used to interpolate between the FFT grid points. The Fourier interpolation provides a very accurate solution for a continuous medium, nevertheless, as it can be seen in Fig. 11, the

Gibbs oscillations are important for a discontinuous medium. To mitigate this problem, we have tried to interpolate linearly between points of the FFT grid. The linear interpolation removes the Gibbs oscillations, but a significant difference with the finite element solution is still observed. Nevertheless, for many applications, such a result would be accurate enough. Note that we have tried different small scale filters applied to the elastic tensor to remove this Gibbs oscillation. Such filters achieve their objective to eliminate the Gibbs oscillations, but as they correspond to a naive upscaling, they also affect the effective elastic tensor, leading to a phase shift in the homogenized seismogram. Such solutions should therefore be avoided. The only accurate solution in that direction would be to use the finite element homogenization at small scales to remove the discontinuities and then use the FFH algorithm in a nested manner. We will investigate such a solution for the 3-D case in a future work.

6. Discussion and conclusion

In this work, we have proposed to apply a FFT iterative algorithm, developed by [11], to solve the cell problem embedded in non-periodic homogenization for the elastic waves. We name it the Fast Fourier Homogenization (FFH) method for elastic wave propagation in deterministic non-periodic media. After [11], other improved algorithms have been proposed to obtain a faster convergence [23, 24]. For most geological media, local velocity contrasts typically vary between 10% to 100% (and 100% is rather exceptional), and our experience has shown that, in such media, the original simple algorithm [11] converges within 30 iterations and often within 10 iterations (for a 10^{-4} accuracy convergence test). With such a low number of iterations

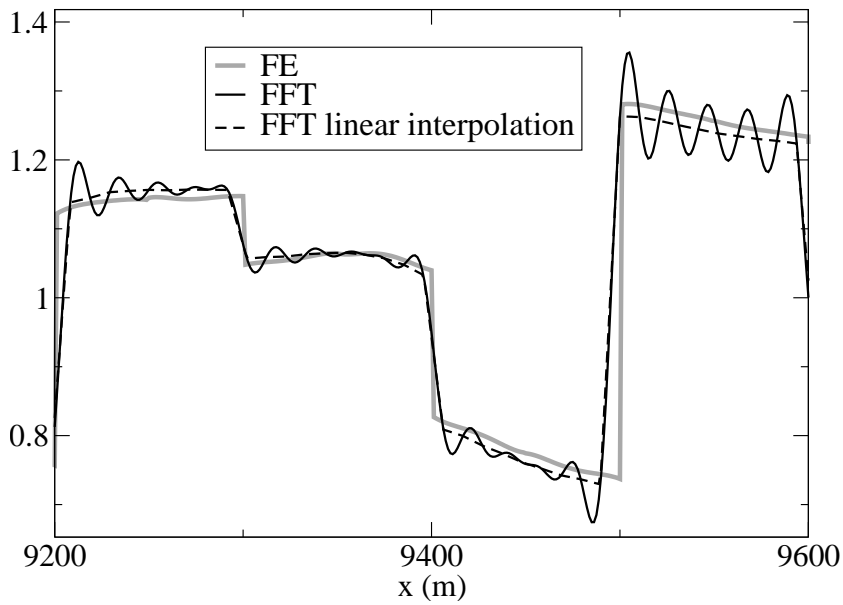


Figure 11: Sample of an horizontal cross-section along $y = 11150 \text{ m}$ in $G_{1111}^{\epsilon_0}$, computed with $\epsilon_0 = 0.25$, with the finite element solver (FE) and with the FFH algorithm with a $dx = 10\text{m}$ sampling and interpolating with the Fourier basis (FFT) and linearly (FFT linear interpolation).

and knowing that the iterative algorithm usually represents only less than half of the total computing time (the rest being input-output, interpolation, etc), an improved iterative algorithm is not really necessary. Nevertheless, knowing that the algorithm developed by [24] is able to handle void and fluid inclusions, it might be useful to implement it in the future to study the effect of void and fluid inclusions on effective properties in the context of wave propagation in non-periodic media.

We have shown that the FFH algorithm is very efficient and accurate when applied to continuous media. It can also be applied to discontinuous media, being a good alternative to finite element homogenization scheme when the mesh is difficult to design and in cases where a regular FFT grid can sample the medium efficiently. At least for the tested cases, it appears that both methods give very similar results for the final homogenization wave propagation, but FFH is much faster than the finite elements, which is a similar conclusion to the one obtained by [22] in the structural mechanics context. Nevertheless, for the correctors, especially for the strain corrector, which is important for sources embedded in heterogeneous media [1], the FFH method must be handled with some care because of the Gibbs phenomenon.

We have shown how the algorithm can be embarrassingly parallelized and how the nested homogenization can solve the buffer issue of such a scheme. This type of parallelization is especially useful for large domains and for users with limited resources.

In conclusion, the FFT iterative homogenization scheme applied to non-periodic homogenization, thanks to its simplicity, its good parallelization capability and efficiency, is a very promising tool. Its 3-D capability will be

presented in a future work. The FFH program is available upon request.

7. Acknowledgments

The authors thank O. Castelnau who, while giving a seminar in Nantes, gave us the idea to use this fantastic algorithm in our scientific context. We also thank two anonymous reviewers and William Parnell who greatly help to improve the manuscript. This work was funded by the the Agence National de la Recherche (ANR) “mémé” project under the blanc program (ANR-10-BLANC-613 MEME). Computations were performed on the CCIPL Erdre cluster. Ming Zhao’s work is partly funded by a Pays de la Loire regional grant and by the Chinese Academy of Sciences. We thank Gaetano Festa for letting us use and modify his 2-D spectral element program.

References

- [1] Y. Capdeville, L. Guillot, J. J. Marigo, 2D nonperiodic homogenization to upscale elastic media for P-SV waves, *Geophys. J. Int.* 182 (2010) 903–922.
- [2] L. Guillot, Y. Capdeville, J. J. Marigo, 2-D non periodic homogenization for the SH wave equation, *Geophys. J. Int.* 182 (2010) 1438–1454.
- [3] Y. Capdeville, L. Guillot, J. J. Marigo, 1-D non periodic homogenization for the wave equation, *Geophys. J. Int.* 181 (2010) 897–910.
- [4] E. Sanchez-Palencia, Non homogeneous media and vibration theory, no. 127 in *Lecture Notes in Physics*, Springer, Berlin, 1980.

- [5] G. Backus, Long-wave elastic anisotropy produced by horizontal layering, *J. Geophys. Res.* 67 (11) (1962) 4427–4440.
- [6] Y. Capdeville, É. Stutzmann, N. Wang, J.-P. Montagner, Residual homogenization for seismic forward and inverse problems in layered media, *Geophys. J. Int.* 194 (1) (2013) 470–487.
- [7] Y. Capdeville, J. J. Marigo, Shallow layer correction for spectral element like methods, *Geophys. J. Int.* 172 (2008) 1135–1150.
- [8] A. Fichtner, H. Igel, Efficient numerical surface wave propagation through the optimization of discrete crustal models- a technique based on non-linear dispersion curve matching (DCM), *Geophys. J. Int.* 173 (2) (2008) 519–533.
- [9] V. Lekić, M. Panning, B. Romanowicz, A simple method for improving crustal corrections in waveform tomography, *Geophys. J. Int.* 182 (1) (2010) 265–278.
- [10] A. Fichtner, J. Trampert, P. Cupillard, E. Saygin, T. Taymaz, Y. Capdeville, A. Villaseñor, Multiscale full waveform inversion, *Geophys. J. Int.* 194 (1) (2013) 534–556.
- [11] H. Moulinec, P. Suquet, A numerical method for computing the overall response of nonlinear composites with complex microstructure, *Comput. Methods Appl. Mech. Engrg.* 157 (1) (1998) 69–94.
- [12] E. Weinan, B. Engquist, X. Li, W. Ren, E. Vanden-Eijnden, Heterogeneous multiscale methods: a review, *Commun. Comput. Phys* 2 (3) (2007) 367–450.

- [13] Y. Capdeville, J.-J. Marigo, A non-periodic two scale asymptotic method to take account of rough topographies for 2-D elastic wave propagation, *Geophys. J. Int.* 192 (1) (2013) 163–189.
- [14] Y. Capdeville, J. J. Marigo, Second order homogenization of the elastic wave equation for non-periodic layered media, *Geophys. J. Int.* 170 (2007) 823–838.
- [15] J. T. Browaeys, S. Chevrot, Decomposition of the elastic tensor and geophysical applications, *Geophys. J. Int.* 159 (2004) 667–678.
- [16] M. Frigo, S. G. Johnson, The design and implementation of FFTW3, *Proceedings of the IEEE* 93 (2) (2005) 216–231, special issue on “Program Generation, Optimization, and Platform Adaptation”.
- [17] H. Dumontet, Study of a boundary layer problem in elastic composite materials, *RAIRO Modél. Math. Anal. Numér.* 20 (2) (1986) 265–286.
- [18] H. Dumontet, Homogénéisation et effets de bords dans les matériaux composites, Ph.D. thesis, Université Paris 6 (1990).
- [19] D. Komatitsch, J. P. Vilotte, The spectral element method: an effective tool to simulate the seismic response of 2D and 3D geological structures, *Bull. Seism. Soc. Am.* 88 (1998) 368–392.
- [20] E. Chaljub, D. Komatitsch, Y. Capdeville, J.-P. Vilotte, B. Valette, G. Festa, Spectral element analysis in seismology, in: R.-S. Wu, V. Maupin (Eds.), *Advances in Wave Propagation in Heterogeneous Media*, Vol. 48 of *Advances in Geophysics Series*, Elsevier, 2007, pp. 365–419.

- [21] G. Festa, J.-P. Vilotte, The newmark scheme as velocity-stress time-staggering: an efficient implementation for spectral element simulations of elastodynamics, *Geophys. J. Int.* 161 (2005) 789–812.
- [22] J. Michel, H. Moulinec, P. Suquet, Effective properties of composite materials with periodic microstructure: a computational approach, *Comput. methods Appl. Mech. Engrg.* 172 (1) (1999) 109–143.
- [23] D. Eyre, G. Milton, A fast numerical scheme for computing the response of composites using grid refinement, *Eur. Phys. J. Appl. Phys.* 6 (1999) 41–48.
- [24] J. Michel, H. Moulinec, P. Suquet, A computational scheme for linear and non-linear composites with arbitrary phase contrast, *Int. J. Numer. Meth. Engrg.* 52 (1-2) (2001) 139–160.

Appendix A. The nested homogenization algorithm

To formalize the nested homogenization idea, we set up a series of values for ε_0 such that

$$\varepsilon_0 > \varepsilon_{0,1} > \dots > \varepsilon_{0,i} > \dots > \varepsilon_{0,N} \quad (\text{A.1})$$

corresponding to a set of wave-numbers

$$k_0 < k_{0,1} < \dots < k_{0,i} < \dots < k_{0,N} \quad (\text{A.2})$$

such that, for each i ,

$$(w^{k_{0,i+1}} * w^{k_{0,i}})(\mathbf{x}) = w^{k_{0,i}}(\mathbf{x}), \quad (\text{A.3})$$

or, in the spectral domain,

$$(\hat{w}^{k_{0,i+1}} \hat{w}^{k_{0,i}})(\mathbf{k}) = \hat{w}^{k_{0,i}}(\mathbf{k}). \quad (\text{A.4})$$

An example of such nested wavelet spectra is given in Fig. A.12.

Based on this series, we build a series of homogenized elastic tensor, defined recursively,

$$\begin{aligned} \mathbf{c}^N &= \mathcal{H}^{k_{0,N}}(\mathbf{c}) \\ \mathbf{c}^{i-1} &= \mathcal{H}^{k_{0,i-1}}(\mathbf{c}^i) \text{ for } i \in \{N..1\}, \end{aligned} \quad (\text{A.5})$$

where $\mathbf{c}^{i=0} = \mathbf{c}^{\varepsilon_0}$ is the desired result.

To obtain the final correctors, we need to combine them together. Using the first homogenization (in ε_0^N) of the nested algorithm, we obtain:

$$\mathbf{u} = \mathbf{u}^{\varepsilon_{0,N}} + \varepsilon_{0,N} \boldsymbol{\chi}^{\varepsilon_{0,N}} : \boldsymbol{\epsilon}(\mathbf{u}^{\varepsilon_{0,N}}) + O(\varepsilon_{0,N}). \quad (\text{A.6})$$

Applying the homogenization in $\varepsilon_{0,N-1}$ to $\mathbf{u}^{\varepsilon_{0,N}}$, we have

$$\mathbf{u}^{\varepsilon_{0,N}} = \mathbf{u}^{\varepsilon_{0,N-1}} + \varepsilon_{0,N-1} \boldsymbol{\chi}^{\varepsilon_{0,N-1}} : \boldsymbol{\epsilon}(\mathbf{u}^{\varepsilon_{0,N-1}}) + O(\varepsilon_{0,N-1}). \quad (\text{A.7})$$




Figure A.12: Examples of three ($N = 2$) nested wavelet $\hat{w}^{k_0,i}(k)$, for $i \in \{0, 1, 2\}$ in the spectral domain.

Combining the last two equations, we obtain, to the first order in $\varepsilon_{0,N-1}$,

$$\mathbf{u} = \mathbf{u}^{\varepsilon_{0,N-1}} + \varepsilon_{0,N-1} \boldsymbol{\chi}^{\varepsilon_{0,N-1}} : \boldsymbol{\epsilon}(\mathbf{u}^{\varepsilon_{0,N-1}}) + \varepsilon_{0,N} \boldsymbol{\chi}^{\varepsilon_{0,N}} : \boldsymbol{\epsilon}(\mathbf{u}^{\varepsilon_{0,N-1}}) + O(\varepsilon_{0,N-1}). \quad (\text{A.8})$$

Repeating the operation until ε_0 is reached, the nested algorithm leads to:

$$\mathbf{u} = \mathbf{u}^{\varepsilon_0} + \sum_{i=1,N} \varepsilon_{0,i} \boldsymbol{\chi}^{\varepsilon_{0,i}} : \boldsymbol{\epsilon}(\mathbf{u}^{\varepsilon_0}) + O(\varepsilon_0). \quad (\text{A.9})$$

Similarly, we have

$$\boldsymbol{\sigma} = \prod_{i=1,N} \mathbf{H}^{\varepsilon_{0,i}} : \boldsymbol{\epsilon}(\mathbf{u}^{\varepsilon_0}) + O(\varepsilon_0) \quad (\text{A.10})$$

$$\boldsymbol{\epsilon} = \prod_{i=1,N} \mathbf{G}^{\varepsilon_{0,i}} : \boldsymbol{\epsilon}(\mathbf{u}^{\varepsilon_0}) + O(\varepsilon_0). \quad (\text{A.11})$$

The nested homogenization therefore allows to find the effective medium and the correctors. The advantage of such a procedure only appears for

parallel implementation where the buffer around each sub-domain needs to be small to allow a good parallel scaling (see Fig. 3). For a given elastic model, whose Fourier spectrum has a maximum wave-number k_{\max} very large compared to the target k_0 , it is always possible to find a $k_{0,N}$ close enough to k_{\max} such that the buffer size is small compared to the sub-domain for the first step of the nested homogenization step. For the next step of the nested homogenization, the maximum wave-number of the $\mathbf{c}^{\varepsilon_{0,N}}$ to be homogenized is proportional to $k_{0,N}$, smaller than the k_{\max} of the original elastic tensor \mathbf{c} . This allows a sparser sampling for the FFT than from the previous step, therefore allowing larger sub-domains. Then, step by step, the series of $k_{0,i}$ leading to k_0 can be built. The nested homogenization is tested in Sec. 5.