

Effective Response of Heterogeneous Materials using the Recursive Projection Method

Xiaoyao Peng (彭瀟瑤)*, Dhriti Nepal†, Kaushik Dayal‡

*‡Department of Civil and Environmental Engineering, Carnegie Mellon University

†Composites Branch, Air Force Research Laboratory

‡Center for Nonlinear Analysis, Carnegie Mellon University

‡Department of Materials Science and Engineering, Carnegie Mellon University

March 6, 2020

Abstract

This paper applies the Recursive Projection Method (RPM) to the problem of finding the effective mechanical response of a periodic heterogeneous solid. Previous works apply the Fast Fourier Transform (FFT) in combination with various fixed-point methods to solve the problem on the periodic unit cell. These have proven extremely powerful in a range of problems ranging from image-based modeling to dislocation plasticity. However, the fixed-point iterations can converge very slowly, or not at all, if the elastic properties have high contrast, such as in the case of voids. The paper examines the reasons for slow, or lack of convergence, in terms of a variational perspective. In particular, when the material contains regions with zero or very small stiffness, there is lack of uniqueness, and the energy landscape has flat or shallow directions. Therefore, in this work, the fixed-point iteration is replaced by the RPM iteration. The RPM uses the fixed-point iteration to adaptively identify the subspace on which fixed-point iterations are unstable, and performs Newton iterations only on the unstable subspace, while fixed-point iterations are performed on the complementary stable subspace. This combination of efficient fixed-point iterations where possible, and expensive but well-convergent Newton iterations where required, is shown to lead to robust and efficient convergence of the method. In particular, RPM-FFT converges well for a wide range of choices of the reference medium, while usual fixed-point iterations are usually sensitive to this choice.

1 Introduction

This paper proposes and applies the Recursive Projection Method (RPM) – a method to adaptively combine Newton and fixed-point iterations – to the problem of finding the effective mechanical response of a

*Email: xiaoyaop@andrew.cmu.edu

†Email: dhriti.nepal.1@us.af.mil

‡Email: Kaushik.Dayal@gmail.com

periodic heterogeneous linear elastic solid. Specifically, given the average strain or stress, the goal is to compute the stress and strain distribution within the unit cell. While the linear elastic setting is of interest in itself, it is also used extensively as part of the iterative solution process in nonlinear elastic or inelastic settings.

Early numerical works in the periodic setting exploited Fast Fourier Transforms (FFT)¹ to partially simplify the problem; however, due to the heterogeneity of the elastic properties, the problem is not completely decoupled in Fourier space. Therefore, fixed-point methods are used to complete the solution process [MS94, MS98, MMS99]. This approach was used to compute the effective properties of heterogeneous materials that are linear, or nonlinear through iteration. However, while the method is extremely fast for small elastic contrasts, the convergence is slow or even lost for higher elastic contrasts. This rules out systems such as materials with voids, where the void is treated as a part of the body with zero stiffness, for instance. This led to the formulation of the Accelerated FFT Method in [MMS01], based on a scheme proposed in [EM99] for scalar conductivity problems, that modified the original method to handle high contrast composites. A further refinement, proposed in [MB12], is the Polarization FFT Method, which has been shown to be a special case of a general class of methods in [MS14]. A separate promising approach is [PL13], based on ideas from energy minimization, though they do not test it on the case of large elastic contrasts.

These methods have been applied in a variety of settings, including in highly nonlinear problems where they are used in each step of the iteration. For instance, in the context of plasticity, we mention application at the macroscale [Leb01, LR20], to continuous dislocation models [BH16, PMB⁺20], as well as recent seminal work in coupling discrete dislocations with fast calculation of elastic interactions [GRL16]. Another important area of application is to phase-field models of microstructure formation in multifunctional materials, e.g. [CS13]. The Fourier methods have also proven powerful in simulating experimental images of microstructure because of the ability to apply the method directly on image data without needing the identification of material boundaries and related challenges [PMG⁺15].

The methods mentioned above – starting from the original FFT method to the Accelerated and Polarization FFT methods – are based on fixed-point methods or variations thereof; see [MS14] for a summary. Consequently, though some of these methods can converge for large elastic contrasts, the convergence gets slower as the contrast increases. In addition, an important part of these methods is the notion of a homogeneous reference medium (for details, see Section 2). The notion of a reference medium goes back to [Esh59], and has proven to be a powerful concept in homogenization analysis [AP13, Lip93]. However, in the numerical setting, the convergence of the fixed-point methods can be very sensitive to the choice of the reference medium. These reasons motivate the proposed work in applying the RPM to this problem.

A classical result is that the solution of the elasticity problem can be posed as a minimization of the strain energy density under standard assumptions, e.g. [Cia88]. That is, the balance of linear momentum, represented by $\operatorname{div} \boldsymbol{\sigma} = 0$ where $\boldsymbol{\sigma}$ is the stress and \mathbf{y} is the deformation, is equivalent to $\min_{\mathbf{y}} \int_{\Omega} W(\nabla \mathbf{y})$, where W is the elastic energy density, Ω is the domain, and $\frac{\partial W}{\partial \nabla \mathbf{y}} = \boldsymbol{\sigma}$. For simplicity, we have assumed displacement fixed on the entire boundary. The perspective that the solution of the equilibrium equation is equivalent to the minimization of strain energy leads to useful insights into the solution methods. In particular, consider the case of a heterogeneous elastic material with voids. Typically, the voids are not

¹Following common practice in this field, we denote also the class of algorithms that use Fast Fourier Transforms as FFT methods.

taken to be part of the body Ω , and instead they are defined through interior surfaces that are traction-free. However, in the FFT method, one cannot easily treat the voids in this manner because the discretization must be uniform for efficiency; hence, the voids are taken to be part of Ω , and treated as regions with zero stiffness. The zero stiffness regions imply that there is infinite elastic contrast, but also *loss of uniqueness* of the solution.

The uniqueness of solutions is an elementary calculation in linear elasticity, e.g. [Cia88]. Assuming a boundary value problem with 2 solutions \mathbf{y}_1 and \mathbf{y}_2 , it can be shown that the difference $\mathbf{u} := \mathbf{y}_1 - \mathbf{y}_2$ satisfies $\int_{\Omega} \boldsymbol{\epsilon}_{\mathbf{u}} : \mathbf{C} \boldsymbol{\epsilon}_{\mathbf{u}} = 0$, where \mathbf{C} is the linear elastic stiffness tensor, and $\boldsymbol{\epsilon}_{\mathbf{u}}$ is the small strain associated with \mathbf{u} . If \mathbf{C} is pointwise positive-definite, then $\boldsymbol{\epsilon}_{\mathbf{u}} \equiv 0$, and we have uniqueness up to rigid motions; Dirichlet boundary conditions on a finite part of the boundary fixes the rigid motion.

From a heuristic perspective, the lack of uniqueness of the solution in a material with zero-stiffness regions is reasonable: the fictitious material points in the void region can have arbitrary displacement (while respecting the required smoothness), and this does not cost any elastic energy or affect the physics of the problem. Hence, we can expect soft or flat directions in energy landscape. When we have large elastic contrast due to regions with very low stiffness, the energy will be shallow though perhaps not completely flat. Convergence along these directions can be exceedingly slow, or even lost, in a fixed-point method.

We briefly note that for rigid inclusions, the strain within the rigid objects must vanish, while the stress can be arbitrary as long as smoothness and equilibrium are satisfied. The discussion above can be readily adapted to this setting using the *complementary* strain energy density.

Given the expectation that the energy landscape is flat or shallow in some directions, the use of Newton or similar iterations is likely to provide convergence within a relatively small number of iterations. However, each Newton iteration is far more expensive than a fixed-point iteration, and increasingly so as the dimension of the solution space increases. The RPM, proposed in [SK93], provides the possibility of a balance between the fixed-point and the Newton methods. The first key aspect of RPM is that it uses the fixed-point iterations to adaptively identify the unstable subspace. That is, starting from the assumption that fixed-point methods will converge when applied to the problem, RPM examines successive outcomes from the fixed-point iteration scheme to discern if, and how, the basis of the unstable subspace should be augmented. The second key aspect of RPM is that as the unstable subspace is built up, Newton iterations are performed on the unstable subspace. This enables a balance between the less expensive but non-converging fixed-point method and the more expensive but better converging Newton method.

We point out that while Newton iterations are performed only on the unstable subspace, the fixed-point iterations are performed on the *entire* space. The reason to perform fixed-point iterations not only on the stable subspace, but rather on the entire subspace which is slightly more expensive, is that it enables the use of existing fixed-point methods with minimal changes to existing code and algorithm. In fact, we exploit this feature to apply RPM to both the original FFT method [MS98] as well as to the Accelerated FFT method [MS14].

Briefly, the key outcomes of our comparisons of RPM-FFT with the existing methods is that the RPM-FFT method is faster than the other methods for large elastic contrasts. However, an even more important advantage appears to be that the convergence of RPM-FFT is very robust with respect to the choice of the reference homogeneous material, whereas the fixed-point methods are sensitive to this choice; this is particularly important since there are only some heuristic ideas about how to optimally choose the reference material.

For brevity, we refer below to the method proposed in this paper as *RPM-FFT*; the conventional algorithm from [MS98] as *Classical FFT*; the method proposed in [MMS01] as *Accelerated FFT*; and the method proposed in [MB12] as *Polarization FFT*.

The manuscript is organized as follows:

- Section 2 summarizes the Classical FFT method;
- Section 3 describes the RPM-FFT method;
- Section 4 provides a variational perspective;
- Section 5 provides numerical examples and comparisons.

2 Classical FFT Method Based on Fixed-Point Iteration

The original FFT method proposed by Moulinec and Suquet [MS98], and some subsequent improvements, are briefly summarized below in the setting of linear elasticity. The goal is to solve the linear elasticity problem on a unit cell V in a periodic geometry with a heterogeneous medium. The average symmetric strain tensor \mathbf{E} is given.

We use $\mathcal{F}[\cdot], \hat{\cdot}, \mathcal{F}^{-1}[\cdot]$ respectively to denote the Fourier transform, Fourier space representation, and inverse Fourier transform.

The displacement field $\mathbf{u}(\mathbf{x})$ is decomposed into a linear part and a zero-mean fluctuating part: $\mathbf{u}(\mathbf{x}) = \mathbf{E} \cdot \mathbf{x} + \mathbf{v}(\mathbf{x})$, ignoring a constant. This implies the strain decomposition $\epsilon_u(\mathbf{x}) = \mathbf{E} + \epsilon_v(\mathbf{x})$, where $\epsilon_{(\cdot)} := \frac{1}{2} (\nabla(\cdot) + \nabla(\cdot)^T)$ is the strain.

We notice that $\hat{\epsilon}_v(\mathbf{k}) = \hat{\epsilon}_u(\mathbf{k}) \forall \mathbf{k} \neq 0$, and $\hat{\epsilon}_u(\mathbf{0}) = \mathbf{E}$ while $\hat{\epsilon}_v(\mathbf{0}) = \mathbf{0}$. Therefore, for conciseness after this section, we will use $\hat{\epsilon}$ to represent both $\hat{\epsilon}_v$ and $\hat{\epsilon}_u$, and ϵ to represent both ϵ_u and ϵ_v .

The unit cell problem can then be written as

$$\operatorname{div} \boldsymbol{\sigma}(\mathbf{x}) = \mathbf{0} \quad \forall \mathbf{x} \in V \quad (2.1)$$

$$\boldsymbol{\sigma}(\mathbf{x}) = \mathbf{C}(\mathbf{x}) \epsilon_u(\mathbf{x}) \quad (2.2)$$

where $\boldsymbol{\sigma}$ is the stress field, and is periodic.

We can then introduce, following [MB12, MS14, EM99] and others, a homogeneous linear elastic comparison medium with stiffness \mathbf{C}^0 and the polarization $\boldsymbol{\tau}(\mathbf{x}) := (\mathbf{C}(\mathbf{x}) - \mathbf{C}^0) \epsilon_u(\mathbf{x})$. This enables us to rewrite the stress-strain relation as:

$$\boldsymbol{\sigma}(\mathbf{x}) = \mathbf{C}^0 \epsilon_v(\mathbf{x}) + \boldsymbol{\tau}(\mathbf{x}) + \mathbf{C}^0 \mathbf{E} \quad (2.3)$$

The solution to the unit cell problem can be written in terms of the periodic Greens function $\boldsymbol{\Gamma}^0$ corresponding to \mathbf{C}^0 :

$$\epsilon_v(\mathbf{x}) = -\boldsymbol{\Gamma}^0(\mathbf{x}) * \boldsymbol{\tau}(\mathbf{x}) \quad \forall \mathbf{x} \in V \quad (2.4)$$

The Fourier space representation of this relation is:

$$\hat{\epsilon}_v(\mathbf{k}) = -\hat{\Gamma}^0(\mathbf{k})\hat{\tau}(\mathbf{k}) \quad \forall \mathbf{k} \neq \mathbf{0}; \hat{\epsilon}_v(\mathbf{0}) = \mathbf{0} \quad (2.5)$$

Typical FFT methods use the fact that taking the (fast) Fourier transform and then multiplying in Fourier space is much faster than convolution in real-space. However, we note that both the real and Fourier representations in (2.4, 2.5) are implicit, through the dependence of the effective forcing τ on the solution \mathbf{u} . Therefore, further work is required to solve this (linear) implicit equation.

From a linear algebra perspective, the Fourier transform takes us to the eigenbasis in a homogeneous problem; in a heterogeneous problem, the Fourier transform brings us close to the eigenbasis, and further calculations are required to completely diagonalize the operator. Roughly, one can make an analogy between the effort to completely diagonalize and the effort to solve the implicit equation in Fourier space.

The implicit equations are typically solved with fixed point methods, and variations of these. In [MS98], for instance, they use the iteration $\hat{\epsilon}^{i+1}(\mathbf{k}) = \hat{\epsilon}^i(\mathbf{k}) - \hat{\Gamma}^0(\mathbf{k})\hat{\sigma}^i(\mathbf{k})$; see Algorithm 1. Like all fixed point methods, convergence is slow or lost when the problem contains an unstable subspace. This is precisely the case when there is large contrast leading to flat energy landscapes and loss of uniqueness, as has been observed in practice [EM99].

Algorithm 1 Classical FFT algorithm [MS98]

```

1: Initialization
2:  $\epsilon^0 \leftarrow \mathbf{E}; \sigma^0 \leftarrow \mathbf{C}(\mathbf{x}) : \epsilon^0$ 
3: Iteration ( $\epsilon^0, \sigma^0$ )
4: while error > tolerance do
5:    $\hat{\epsilon}^i \leftarrow \mathcal{F}[\epsilon^i]; \hat{\sigma}^i \leftarrow \mathcal{F}[\sigma^i]$  ▷ FFT on  $\epsilon$  and  $\sigma$ 
6:    $\hat{\epsilon}^{i+1} \leftarrow \hat{\epsilon}^i - \hat{\Gamma}^0 : \hat{\sigma}^i$  ▷  $\forall \mathbf{k} \neq \mathbf{0}; \hat{\epsilon}^{i+1}(\mathbf{0}) = \mathbf{E}$ 
7:    $\epsilon^{i+1} \leftarrow \mathcal{F}^{-1}[\hat{\epsilon}^{i+1}]$ 
8:    $\sigma^{i+1} \leftarrow \mathbf{C}(\mathbf{x}) : \epsilon^{i+1}$ 
9:   error  $\leftarrow \frac{(\langle \|\mathbf{k} \cdot \hat{\sigma}(\mathbf{k})\|^2 \rangle)^{1/2}}{\|\hat{\sigma}(\mathbf{0})\|} \equiv \frac{(\langle \|\operatorname{div} \sigma\|^2 \rangle)^{1/2}}{\|\langle \sigma \rangle\|}$ 
10: end while

```

3 The Proposed RPM-FFT Method

We first outline the key elements of the recursive projection method following the seminal work of Shroff and Keller [SK93], and then the application of RPM to the FFT approach.

When using fixed-point iterative procedures to solve nonlinear problems, slow convergence or even divergence can be caused by a small number of eigenvalues of the Jacobian matrix approaching or leaving the unit disk $\{|z| < 1\}$. The Recursive Projection Method (RPM) is a stabilization procedure that can recover the convergence.

Consider a problem in N dimensions posed on \mathbb{R}^N . The method begins by dividing the space \mathbb{R}^N into \mathbb{P} , that we denote as the *unstable eigenspace*, and its orthogonal complement \mathbb{Q} . \mathbb{P} and \mathbb{Q} correspond

respectively to the eigenspace of those eigenvalues leaving the disk and those inside the disk. Then, the fixed-point iteration is performed only on the complement \mathbb{Q} while Newton's method – or, in practice, a variation such as quasi-Newton – is performed on the unstable eigenspace \mathbb{P} .

Writing the original problem to be solved in the fixed-point form as follows:

$$\mathbf{u} = \mathbf{F}(\mathbf{u}), \quad \mathbf{F}: \mathbb{R}^N \rightarrow \mathbb{R}^N \quad (3.1)$$

the fixed-point iteration can be written:

$$\mathbf{u}^{(\nu+1)} = \mathbf{F}(\mathbf{u}^{(\nu)}) \quad (3.2)$$

where ν indicates the iteration number. While the eigenvalues of the Jacobian matrix $\mathbf{F}_u := \frac{\partial \mathbf{F}}{\partial \mathbf{u}}$ lie inside of the disk

$$K_\delta = \{|z| \leq 1 - \delta\} \quad \text{for some } \delta > 0 \quad (3.3)$$

and the initial value $\mathbf{u}^{(0)}$ is sufficiently close to the solution, the fixed-point iteration will converge. A small positive constant δ is used to ensure that one is sufficiently far from the boundary between stable and unstable fixed point iterations.

Suppose now that m eigenvalues lie outside the disk K_δ

$$|\mu_1| \geq \cdots \geq |\mu_m| > 1 - \delta \geq |\mu_{m+1}| \geq \cdots \geq |\mu_N| \quad (3.4)$$

Let $\mathbf{Z} \in \mathbb{R}^{N \times m}$ be the orthonormal basis for \mathbb{P} . The projectors of the subspaces \mathbb{P} and \mathbb{Q} are respectively $\mathbf{P} = \mathbf{Z}\mathbf{Z}^T$ and $\mathbf{Q} = \mathbf{I} - \mathbf{Z}\mathbf{Z}^T$. For each $\mathbf{u} \in \mathbb{R}^n$, there is a unique decomposition:

$$\mathbf{u} = \mathbf{p} + \mathbf{q}, \quad \mathbf{p} \equiv \mathbf{P}\mathbf{u} \in \mathbb{P}, \quad \mathbf{q} \equiv \mathbf{Q}\mathbf{u} \in \mathbb{Q} \quad (3.5)$$

This allows the introduction of $\mathbf{f}(\mathbf{p}, \mathbf{q}) := \mathbf{P}\mathbf{F}(\mathbf{p} + \mathbf{q})$ and $\mathbf{g}(\mathbf{p}, \mathbf{q}) := \mathbf{Q}\mathbf{F}(\mathbf{p} + \mathbf{q})$.

Applying Newton's method on the unstable subspace \mathbb{P} :

$$\mathbf{p}^{(\nu+1)} = \mathbf{p}^{(\nu)} + (\mathbf{I} - \mathbf{f}_p^{(\nu)})^{-1}(\mathbf{f}(\mathbf{p}^{(\nu)}, \mathbf{q}^{(\nu)}) - \mathbf{p}^{(\nu)}) \quad (3.6)$$

where $\mathbf{f}_p(\mathbf{p}^{(\nu)}, \mathbf{q}^{(\nu)}) = \mathbf{P}\mathbf{F}_u(\mathbf{u}^{(\nu)})\mathbf{P}$ is the derivative of \mathbf{f} with respect to \mathbf{p} .

We continue to use fixed-point iteration on \mathbb{Q} :

$$\mathbf{q}^{(\nu+1)} = \mathbf{g}(\mathbf{p}^{(\nu)}, \mathbf{q}^{(\nu)}) \quad (3.7)$$

Efficiently identifying the unstable eigenspace \mathbb{P} is an essential feature of RPM. This corresponds to finding the basis \mathbf{Z} ; or, in practice, an approximation of \mathbf{Z} since the Jacobian matrix can be expensive to compute. In [SK93], this is accomplished directly by monitoring the convergence rate of $\mathbf{q}^{(\nu)}$, without computing the Jacobian matrix. This is an important feature that significantly reduces the computational cost, and we follow that procedure here.

The method begins by assuming that the unstable eigenspace \mathbb{P} is 0-dimensional and $\mathbf{Z} = \mathbf{0}$. We then iteratively build up the unstable eigenspace by adding to \mathbf{Z} the eigenspace corresponding to those eigenvalues that approach the unit disk. The eigenspace that corresponds to the eigenvalues approaching the unit disk corresponds to the dominant eigenspace of $\mathbf{g}_q = \mathbf{Q}\mathbf{F}_u\mathbf{Q}$.

If the fixed-point iteration does not converge within the specified limit of iterations n_{max} , the eigenspace corresponding to the eigenvalues approaching the unit disk is identified as follows. In [SK93], it is shown that $\Delta\mathbf{q}^{(\nu)} \equiv \mathbf{q}^{(\nu+1)} - \mathbf{q}^{(\nu)} \approx (\mathbf{g}_q)^\nu \Delta\mathbf{q}^{(0)}$, i.e., a power iteration with \mathbf{g}_q applied to $\Delta\mathbf{q}^{(0)}$. From the properties of power iterations, $\Delta\mathbf{q}^{(\nu)}$ identifies the eigenspace of \mathbf{g}_q that corresponds to the largest eigenvalue, assuming that $\Delta\mathbf{q}^{(0)}$ has a nonzero component in this space.

Following [SK93], we approximate the dominant eigenspace of \mathbf{g}_q using QR factorization to write:

$$\mathbf{D} \equiv [\Delta\mathbf{q}^{(\nu)}, \Delta\mathbf{q}^{(\nu-1)}] =: \tilde{\mathbf{D}}\mathbf{T} \quad (3.8)$$

where $\tilde{\mathbf{D}} \in \mathbb{R}^{N \times 2}$ is orthogonal and $\mathbf{T} \in \mathbb{R}^{2 \times 2}$ is upper triangular. When $\mathbf{T}_{11} \gg \mathbf{T}_{22}$, we add the first column of $\tilde{\mathbf{D}}$ to the basis \mathbf{Z} , corresponding to one real eigenvalue approaching the unit disk. Else, we add both columns of $\tilde{\mathbf{D}}$ to the basis \mathbf{Z} , corresponding to a complex conjugate pair of eigenvalues approaching the unit disk. Alternate approximations are further discussed in [SK93].

To apply the RPM method to the specific context of the FFT fixed-point algorithm, we formally define the fixed-point operator \mathbf{F} from Algorithm 1 through the equation:

$$\boldsymbol{\epsilon}^{i+1} = \mathcal{F}^{-1} \left[\mathcal{F}[\boldsymbol{\epsilon}^i] - \hat{\mathbf{I}}^0 : \mathcal{F}[\mathbf{C}(\mathbf{x}) : \boldsymbol{\epsilon}^i] \right] \quad (3.9)$$

Define $\mathbf{z} := \mathbf{Z}^T \mathbf{p} \in \mathbb{R}^m$ to represent \mathbf{p} in the basis \mathbf{Z} . Then the stable / unstable decomposition can also be written as:

$$\mathbf{u} = \mathbf{Z}\mathbf{z} + \mathbf{q}, \quad \mathbf{p} = \mathbf{Z}\mathbf{z}, \quad \mathbf{q} = (\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\mathbf{u} \quad (3.10)$$

It is suggested in [SK93] that it is often sufficient for convergence to simply compute only once the quantity $(\mathbf{I} - \mathbf{f}_p)^{-1}$ using finite differences. In our calculations in this paper, we have also found this to be sufficient for convergence, and of course very efficient.

Defining $\mathbf{H} := \mathbf{Z}^T \mathbf{F}_u \mathbf{Z}$, the iteration (3.6, 3.7) can now be rewritten in a form that is more efficient and transparent for implementation:

$$\mathbf{z}^{(\nu+1)} = \mathbf{z}^{(\nu)} + (\mathbf{I} - \mathbf{H})^{-1}(\mathbf{Z}^T \mathbf{F}(\mathbf{u}^{(\nu)}) - \mathbf{z}^{(\nu)}) \quad (3.11)$$

$$\mathbf{q}^{(\nu+1)} = (\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)\mathbf{F}(\mathbf{u}^{(\nu)}) \quad (3.12)$$

$$\mathbf{u}^{(\nu+1)} = \mathbf{Z}\mathbf{z}^{(\nu+1)} + \mathbf{q}^{(\nu+1)} \quad (3.13)$$

Finally, integrating the FFT method into RPM, gives us Algorithm 2.

An important feature is that the fixed-point iteration is conducted on the entire solution space and not only on the stable subspace. After the fixed-point iteration, only that part of the outcome corresponding to the stable part is retained in constructing the next iterate of the approximate solution. While marginally more computationally expensive than using a fixed-point iteration only on the stable subspace, it has the important advantage that standard existing fixed-point codes can directly be used with minimal modification as the RPM “wraps around” the fixed-point method.

4 A Variational Perspective

We begin with noting an analogy between the problem of interest here and the energy-minimization formulation of electrostatic fields in matter. The field equations and constitutive response of electrostatic

Algorithm 2 RPM-FFT algorithm

```

1: Initialization
2:  $\epsilon^0 \leftarrow \mathbf{E}$ ;  $\xi^0 \leftarrow \mathbf{F}(\epsilon^0)$ 
3: Iteration ( $\epsilon^0, \xi^0$ )
4: while error > tolerance do
5:    $\mathbf{z}^{(\nu)} \leftarrow \mathbf{Z}^T \epsilon^{(\nu)}$ ;  $\zeta^{(\nu)} \leftarrow \mathbf{Z}^T \xi^{(\nu)}$                                  $\triangleright$  Project to unstable subspace
6:    $\mathbf{z}^{(\nu+1)} \leftarrow \mathbf{z}^{(\nu)} + (\mathbf{I} - \mathbf{H})^{-1}(\zeta^{(\nu)} - \mathbf{z}^{(\nu)})$            $\triangleright$  Newton iteration
7:    $\mathbf{q}^{(\nu+1)} \leftarrow \xi^{(\nu)} - \mathbf{Z} \zeta^{(\nu)}$ 
8:    $\epsilon^{(\nu+1)} \leftarrow \mathbf{Z} \mathbf{z}^{(\nu+1)} + \mathbf{q}^{(\nu+1)}$ 
9:    $\xi^{(\nu+1)} \leftarrow \mathbf{F}(\epsilon^{(\nu+1)})$                                                $\triangleright$  Fixed-point iteration
10:   $\nu = \nu + 1$ 
11:  if  $\nu > n_{max}$  then       $\triangleright$  Start Newton iteration if number of fixed-point iterations exceeds  $n_{max}$ 
12:     $\mathbf{Z} \leftarrow \text{In}(\mathbf{q}, \mathbf{Z})$                                  $\triangleright$  Increase basis size
13:     $\mathbf{H} \leftarrow \mathbf{F}_u \mathbf{Z}$                                  $\triangleright$  Compute additional column of  $\mathbf{H}$ 
14:     $\nu = 0$ 
15:  end if
16:  error  $\leftarrow \frac{(\langle \|\mathbf{k} \cdot \hat{\sigma}(\mathbf{k})\|^2 \rangle)^{1/2}}{\|\hat{\sigma}(\mathbf{0})\|}$ 
17: end while

```

fields in matter are respectively:

$$\text{div } \nabla \phi = \text{div } \mathbf{p} \text{ on } \Omega, \quad \mathbf{p}(\mathbf{x}) = \tilde{\mathbf{p}}(\nabla \phi(\mathbf{x}), \mathbf{x}), \quad \text{BCs: } \phi(\mathbf{x}) = V_0(\mathbf{x}) \quad \forall \mathbf{x} \in \partial \Omega \quad (4.1)$$

where ϕ is the electric potential, \mathbf{p} is the dipole per unit volume induced by the electric field $\nabla \phi$, and $\tilde{\mathbf{p}}$ is the material-dependent response function that can be an explicit function of position. For simplicity, we have assumed Dirichlet boundary conditions V_0 on the boundary $\partial \Omega$ of the domain Ω , but it is straightforward to use more general BCs. Note that we neglected various constants such as the permittivity of vacuum.

Under some physically-motivated conditions on $\tilde{\mathbf{p}}$, this can be reformulated as the minimization of electrostatic energy. For instance, we consider the case of a linear dielectric which is analogous to the linear elastic medium considered in this paper, in which case $\tilde{\mathbf{p}}$ takes the form:

$$\tilde{\mathbf{p}}(\mathbf{E}(\mathbf{x}), \mathbf{x}) = \chi(\mathbf{x}) \mathbf{E} \quad (4.2)$$

where $\chi(\mathbf{x})$ is the second-order electrical susceptibility tensor.

Following [JK90] for magnetostatics – see also [SB01] and [YD11] for electrostatics – we can write this as an energy minimization with a nonlocal constraint:

$$\min E_{electro}[\mathbf{p}, \phi], \text{ with } E_{electro}[\mathbf{p}, \phi] = \frac{1}{2} \int_{\Omega} (\mathbf{p} \cdot \chi^{-1} \mathbf{p} + |\nabla \phi|^2), \text{ subject to } \text{div } \nabla \phi = \text{div } \mathbf{p} \text{ on } \Omega \quad (4.3)$$

with BCs $\phi(\mathbf{x}) = V_0(\mathbf{x}) \quad \forall \mathbf{x} \in \partial \Omega$.

From the fundamental physics of electrostatics problems, related to the motion of charges under an electric field, we have χ is positive-definite pointwise. Therefore, from an elementary application of the direct method of the calculus of variations, a unique minimizer exists; loosely, the energy is bounded below and grows quadratically in all directions from the positive-definiteness of χ . The nonlocal constraint, while somewhat nonstandard, can be readily dealt with, e.g. [JK90].

We turn now to linear inhomogeneous elasticity. We recall the balance of linear momentum:

$$\operatorname{div} \mathbf{C}(\mathbf{x}) \boldsymbol{\epsilon}_u = 0 \quad (4.4)$$

posed on the periodic unit cell V with periodic BCs. We decompose $\mathbf{C}(\mathbf{x}) = \mathbf{C}_0 + \mathbf{d}(\mathbf{x})$, where \mathbf{C}_0 is the homogeneous reference medium. This gives:

$$\operatorname{div} \mathbf{C}_0 \boldsymbol{\epsilon}_u = -\operatorname{div} \mathbf{d} \boldsymbol{\epsilon}_u \quad (4.5)$$

Recall from Section 2 that we defined the polarization field $\boldsymbol{\tau} := \mathbf{d} \boldsymbol{\epsilon}_u$. Following closely the electrostatic case, we can pose this as a nonlocally-constrained critical-point problem. Define the energy E :

$$E[\boldsymbol{\tau}, \mathbf{u}] = \frac{1}{2} \int_V (\boldsymbol{\tau} : \mathbf{d}^{-1} \boldsymbol{\tau} + \boldsymbol{\epsilon}_u : \mathbf{C}_0 \boldsymbol{\epsilon}_u), \text{ subject to } \operatorname{div} \mathbf{C}_0 \boldsymbol{\epsilon}_u = -\operatorname{div} \boldsymbol{\tau} \quad (4.6)$$

Use the variation $\boldsymbol{\tau} \rightarrow \boldsymbol{\tau} + \epsilon \boldsymbol{\omega}$ and $\mathbf{u} \rightarrow \mathbf{u} + \epsilon \mathbf{w}$:

$$\delta E = \int_V (\boldsymbol{\omega} : \mathbf{d}^{-1} \boldsymbol{\tau} + \boldsymbol{\epsilon}_u : \mathbf{C}_0 \boldsymbol{\epsilon}_w) \quad (4.7)$$

The variations $\boldsymbol{\omega}$ and \mathbf{w} are not independent; they are constrained to satisfy:

$$\operatorname{div} \mathbf{C}_0 \boldsymbol{\epsilon}_w = -\operatorname{div} \boldsymbol{\omega} \quad (4.8)$$

On both sides of the equation above, do the following: (1) take the inner product with $\boldsymbol{\epsilon}_u$; (2) integrate over V ; (3) use integration-by-parts to move the derivatives. The result of these operations is:

$$\int_V \boldsymbol{\epsilon}_u : \mathbf{C}_0 \boldsymbol{\epsilon}_w = \int_V -\boldsymbol{\omega} : \boldsymbol{\epsilon}_u \quad (4.9)$$

where the boundary terms cancel out due to periodicity.

Using (4.9) in (4.7), and requiring $\delta E = 0$ for arbitrary $\boldsymbol{\omega}$, we have:

$$\mathbf{d}^{-1} \boldsymbol{\tau} - \boldsymbol{\epsilon}_u = 0, \quad \operatorname{div} \mathbf{C}_0 \boldsymbol{\epsilon}_u = -\operatorname{div} \boldsymbol{\tau} \quad (4.10)$$

We highlight here that the tensor $\mathbf{d}(\mathbf{x})$ is generally not positive-definite pointwise, particularly in the important case of high elastic contrasts. For instance, consider a material with voids, where the non-voided material is stable, i.e. $\mathbf{C}(\mathbf{x})$ is positive-definite in the non-voided material, and $\mathbf{C}(\mathbf{x}) = \mathbf{0}$ in the void region. The reference medium is taken to be stable, i.e. \mathbf{C}_0 is positive-definite to enable solution of the constraint equation. Then, $\mathbf{d} = -\mathbf{C}_0$ in the voids, and is not positive-definite pointwise. Therefore, in general, we do not expect existence of minimizers for E .

Loosely, the energy has unstable directions which are not bounded below. Therefore, gradient descent methods – which can be written as fixed-point methods as shown below for this energy formulation and by [Sch17] – do not converge. On the other hand, using Newton iterations along the unstable directions – once these are identified – can greatly improve convergence, which is precisely the role of the Recursive Projection Method.

4.1 Fixed-Point Iterations as Gradient Descent

We consider a gradient / steepest descent approach to minimizing the energy E :

$$E = \frac{1}{2} \int_{\Omega_\#} \left(\boldsymbol{\tau} : \mathbf{d}^{-1} \boldsymbol{\tau} + \frac{1}{2} \boldsymbol{\epsilon}_u : \mathbf{C}_0 \boldsymbol{\epsilon}_u \right), \text{ subject to } \operatorname{div} \mathbf{C}_0 \boldsymbol{\epsilon}_u = \operatorname{div} \boldsymbol{\tau} \quad (4.11)$$

The gradient direction in function space is given by $\mathbf{d}^{-1} \boldsymbol{\tau} - \boldsymbol{\epsilon}_u$, and a gradient flow in the standard L^2 norm is $\dot{\boldsymbol{\tau}} \sim \mathbf{d}^{-1} \boldsymbol{\tau} - \boldsymbol{\epsilon}_u$.

Consider an explicit update scheme where the constraint equation is also updated at each iteration. We obtain the following fixed-point algorithm:

1. given $\boldsymbol{\epsilon}_u^i$
2. $\boldsymbol{\tau}^{i+1} = \boldsymbol{\tau}^i + a(\mathbf{d}^{-1} \boldsymbol{\tau}^i - \boldsymbol{\epsilon}_u^i)$
3. $\operatorname{div} \mathbf{C}_0 \boldsymbol{\epsilon}_u^{i+1} = \operatorname{div} \boldsymbol{\tau}^{i+1}$
4. loop

(4.12)

where a is related to the fictitious timestep and mobility, and the superscript indexes the iterations. Notice that the constraint update (step 3 above) requires solution of a homogeneous periodic linear elastic problem with a given right side. This can be done very efficiently using fast Fourier transforms. Each iteration in the loop above is therefore extremely quick, but the key question is whether the algorithm converges and, if so, how many iterations it takes. Given that E may not possess minimizers, it follows that gradient descent may not converge.

The Classical FFT, summarized in [MS14], can be written in the fixed point form:

1. given $\boldsymbol{\epsilon}_u^i$
2. $\boldsymbol{\tau}^{i+1} = \mathbf{d} \boldsymbol{\epsilon}_u^i$
3. $\operatorname{div} \mathbf{C}_0 \boldsymbol{\epsilon}_u^{i+1} = \operatorname{div} \boldsymbol{\tau}^{i+1}$
4. loop

(4.13)

This can be considered as derived from the explicit update scheme above, except that the update for $\boldsymbol{\tau}$ in step 2 goes directly to the minimum – for a given $\boldsymbol{\epsilon}_u^i$ – using that the energy is quadratic. As in the explicit update scheme, this is very fast for each iteration, but whether it converges, and the rate of convergence, depend on the structure of the energy landscape.

The Polarization FFT proposed by [MB12] can be rewritten in the form:

1. given $\boldsymbol{\epsilon}_u^i$
2. $\boldsymbol{\tau}^{i+1} = \mathbf{d} \boldsymbol{\epsilon}_u^i$
3. $\operatorname{div} \mathbf{C}_0 \boldsymbol{\epsilon}_u^{i+1} = \operatorname{div} \mathbf{C}_0 \mathbf{d}^{-1} \boldsymbol{\tau}^{i+1} + \beta \operatorname{div} \mathbf{C}_0 \mathbf{d}^{-1} \mathbf{C}_0 \boldsymbol{\epsilon}_u^i - \operatorname{div} \mathbf{C}_0 \mathbf{d}^{-1} \mathbf{C}_0 \tilde{\boldsymbol{\epsilon}}_u^i$
- where $\tilde{\boldsymbol{\epsilon}}_u^i$ solves $\operatorname{div} \mathbf{C}_0 \tilde{\boldsymbol{\epsilon}}_u^i = \operatorname{div} (\alpha \boldsymbol{\tau}^i + (\alpha + \beta) \mathbf{C}_0 \boldsymbol{\epsilon}_u^i)$
4. loop

(4.14)

While this form is not as transparent or convenient, it enables us to compare the methods in the context of gradient descent. We notice that in this form, the polarization method can be considered as a relaxation method that mixes in the value from the previous iteration, with α and β being the relaxation parameters.

5 Numerical Comparisons Between Classical FFT, Accelerated FFT, and RPM-FFT

As a canonical problem for many of the calculations below, we follow [MS98] in using a circular stiff fiber embedded in a compliant matrix, at the center of a square unit cell and periodic boundary conditions (Fig. 1). This example is chosen because of the availability of approximate closed-form solutions. We use the following notation: a denotes the fiber radius; b denotes the size of the square unit cell; E_f and $\nu_f = 0.35$ denote respectively the elastic modulus and the Poisson's ratio of the fiber, noting that we assume isotropic elasticity; E_m and $\nu_m = 0.23$ denote respectively the elastic modulus and the Poisson's ratio of the matrix, again assuming isotropic elasticity; and $K = E_f/E_m$ is the elastic contrast, where particular focus will be on large values of K . The average strain is denoted by the components E_{12}, E_{11}, E_{22} . For all tests in this section except Section 5.5, the reference medium is assumed to be isotropic and we choose the elastic modulus $E_o = \frac{E_f + E_m}{2}$ and the Poisson's ratio $\nu_o = \nu_f = \nu_m$.

As a first test, we simply confirm that RPM-FFT converges to the correct solution in Fig. 1. There is no discernible difference² in the solutions obtained from the Classical FFT and RPM-FFT methods. Further, Classical FFT takes less time³ than RPM-FFT. This is unsurprising given that the elastic contrast is small ($K \approx 5.8$). As K increases, the advantages of the RPM-FFT method over the Classical FFT method become clear. We examine this and other issues below, with a more detailed reporting of computational time.

5.1 Time to Convergence vs. Elastic Contrast

We next compare the time taken for convergence across the Classical FFT, Accelerated FFT, and RPM-FFT methods. We continue with the circular fiber in matrix setting, with K going from 10 to 10^5 , and $\nu_m = \nu_f = 0.25$. We use a 256×256 discretization, and apply an average shear strain: $E_{12} = 0.5\%$, $E_{11} = E_{22} = 0$.

Fig. 2 shows the now well-known dramatic improvement between the Classical FFT and the Accelerated FFT methods; [MMS01] show that the rate of convergence in the Classical FFT method goes as the contrast K , while for the Accelerated FFT method the rate of convergence goes as the square root of the contrast K .

Fig. 2 also compares the Accelerated FFT method and RPM-FFT method. When the contrast is below 10^3 , the Accelerated FFT method converges faster. However, as the contrast increases, RPM-FFT is increasingly competitive, and surpasses Accelerated FFT. The rate of convergence of RPM goes as $K^{0.43}$ for this particular class of examples. The exponent appears to be roughly the same in a few other problems that we tested.

We note that the comparisons for all of these methods was conducted by implementing the codes in Matlab with no additional optimization of any of the methods to ensure fair comparisons.

²The tolerance in all calculations reported in this paper is taken to be 10^{-4} .

³All times reported in this paper are wall-clock times.

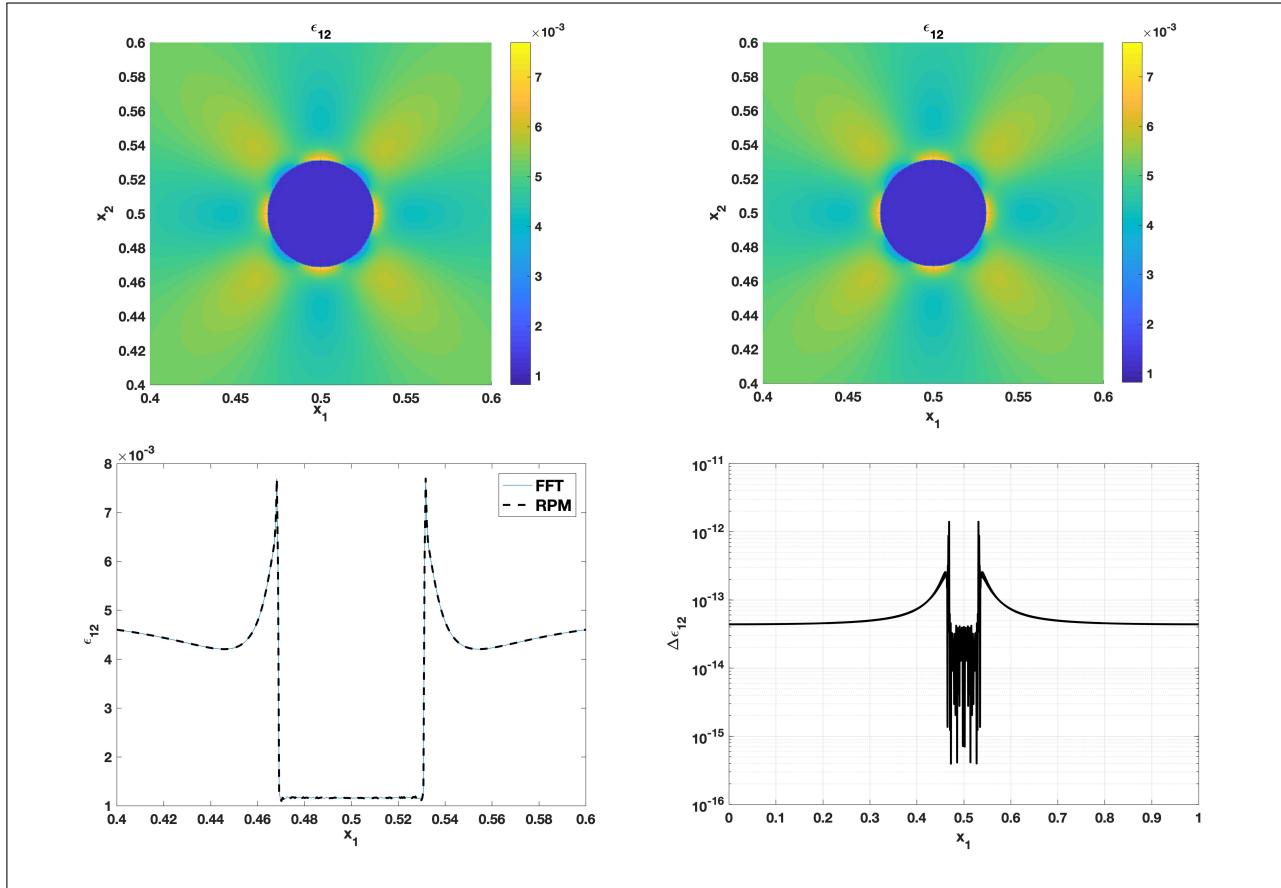


Figure 1: Top row: ϵ_{12} component of the strain using Classical FFT (left) and RPM-FFT (center) methods. Bottom row: the figure on the right shows ϵ_{12} along the horizontal symmetric axis for both RPM-FFT and Classical FFT; the figure on the left shows the normalized difference. The geometry is given by $a/b = 1/16$, and the figure shows only the region in the vicinity of the fiber. The mesh is 2048×2048 . We use $E_f = 68.9$ GPa and $E_m = 400$ GPa; the contrast $K \approx 5.8$. We apply shear loading: $E_{12} = 0.5\%$, $E_{11} = E_{22} = 0$.

5.2 Balancing Newton Iterations against Fixed Point

As can be expected, the speed of RPM-FFT depends strongly on the size of the unstable subspace on which Newton iterations must be performed. While the Newton's method has better convergence properties, each Newton iteration is also far more expensive than the fixed point iteration, particularly as the dimension of the unstable subspace increases. On the other hand, the fixed point method converges slowly, or not at all, on the unstable subspace, which also increases the expense as it requires a large number of iterations.

In this section, the term “unstable subspace” refers to the subspace that the RPM-FFT method identifies; this is not necessarily the true unstable subspace whose corresponding eigenvalues are outside the unit disk.

An optimal method has to balance between the opposing issues identified above. Denoting by n_{max} the number of fixed-point iterations before increasing the size of the unstable subspace, we notice that setting n_{max} too low causes the dimensionality of the unstable subspace to increase quickly and makes

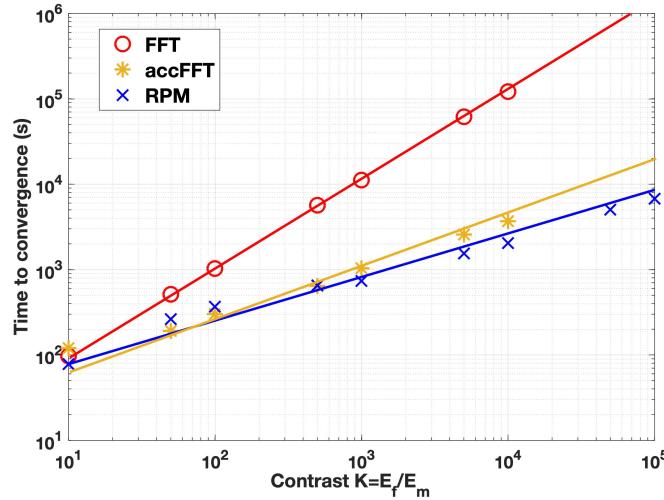


Figure 2: Time at convergence of three methods with varying K . Note the log scale. The lines are linear best fits.

each Newton iteration expensive; on the other hand, if n_{max} is too high, there will be minimal progress from the large number of fixed-point iterations before switching to Newton iterations.

We examine this interplay numerically for the model problem described previously, for varying K . We first fix K , and then compare the times obtained for a large range of n_{max} going from 1 to 50. We then repeat over a range of K . Of these, we denote the value of n_{max} with the best (least) time as $n_{max}^{best}(K)$. Fig. 3 shows the times to convergence for a fixed value of $n_{max} = 10$, as well as the times to convergence for $n_{max}^{best}(K)$, as a function of K . An immediate conclusion is that there is almost no difference at low K , and not much difference at higher K .

We next compare the effect of the size of the problem on $n_{max}^{best}(K)$, for a fixed value of $K = 50$. Fig. 4 shows n_{max}^{best} for different problem sizes, here identified with the number of grid points in each direction. An immediate conclusion is that n_{max}^{best} increases with the problem size.

5.3 Effect of Volume Fraction and Problem Size

In the model problem of a stiff fiber in a soft matrix, the volume fraction of the stiff fiber has been taken to be small; for $a/b = 1/16$, the stiff fiber has a volume fraction $\lambda = 0.3\%$. Note that we use the term volume fraction though we are working in the 2D setting.

We expect that the size of the unstable subspace is related to the size of the fiber based on the heuristics discussed in Sections 1, 3 and 4. As the volume fraction increases, we therefore expect a larger unstable subspace and consequently more expensive Newton iterations.

We examine this issue through numerical experiments, and find that RPM-FFT works well even at the combination of large elastic contrast K and volume fraction of fiber λ . For instance, even with $K \approx 10^4$ or higher, and $a/b = 9/10 \Rightarrow \lambda = 63.6\%$, RPM-FFT is easily able to converge. Fig. 5 shows the time to convergence as a function of λ . While larger elastic contrasts require more time to converge – in line with our observations above – we notice that the time to converge is almost insensitive to λ for volume fractions larger than 5%.

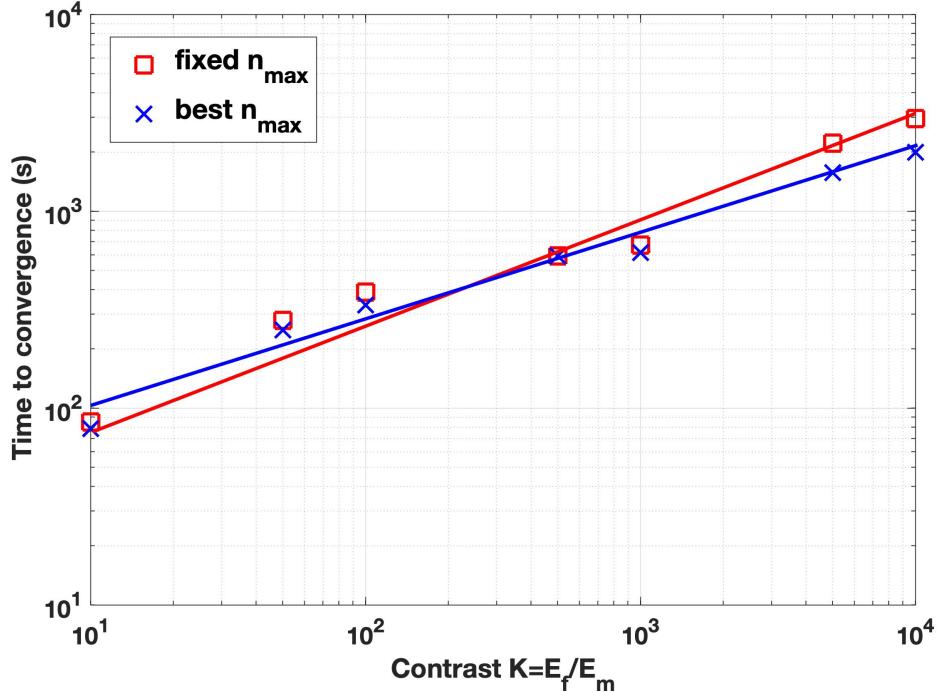


Figure 3: Time at convergence of RPM with fixed $n_{max} = 10$ and $n_{max}^{best}(K)$. The lines are linear best fits.

We next consider the case of fixed λ but with changing problem size with a fixed geometry. That is, we compare discretizations of varying coarseness. The number of grid points within the fiber, as well as overall, increases as the discretization is made finer. We consequently expect the unstable subspace identified by RPM to also increase. Fig. 6 shows the relation between the size of the unstable subspace as the problem size increases from 32×32 to 512×512 , for different values of elastic contrast K . The volume fraction λ is held fixed 1.2%. Roughly, we observe a logarithmic scaling between the size of the unstable subspace and the number of grid points within the fiber.

5.4 Interaction Between Stiff Fibers

Motivated by recent advances in functional composites [HSP⁺19, RRWN17], we examine the setting of two circular carbon fibers embedded in a polyurethane matrix. It is a significant challenge to characterize the effective properties of heterogeneous materials such as nanofiber- and nanoparticle- reinforced nanocomposites starting from the scale of resolving the entire microstructure. This makes it challenging to develop physics-based high-fidelity model that predict the performance of such nanocomposite systems. The FFT methods can help to make stronger direct connections between mechanics modeling and experimental data, which in turn can have a significant impact on understanding the structure-property relations of these complex heterogeneous nanocomposite systems.

We consider the geometry of Fig. 7 with two fibers that are either close to each other – the separation is smaller than the smaller fiber diameter – and when they are further apart. The moduli are assumed to be $E_f = 900$ GPa for the carbon fiber and $E_m = 0.03$ GPa for the polyurethane, giving an elastic contrast of $K = 3 \times 10^4$. Two fibers, radius a and $2a$, are located in a rectangular unit cell of size $2b \times b$, with

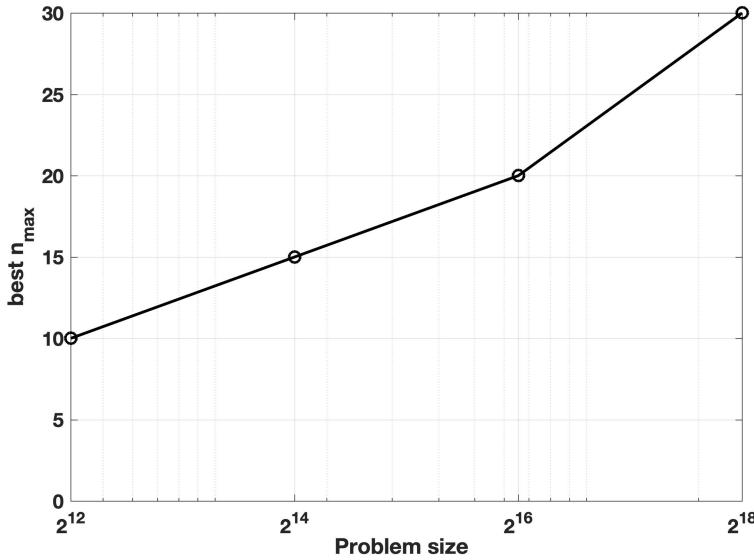


Figure 4: The best n_{max} for different problem sizes. The line connects datapoints to aid visualization.

$b = 10a$. We examine two cases: when the fibers are far apart with distance b between the centers of the fibers, and when the fibers are closer with distance $0.4b$, see Fig. 7. We apply a uniaxial average strain: $E_{11} = 5\%$, $E_{22} = E_{12} = 0$.

Fig. 7 shows strain and strain energy density fields for the two configurations. In all of these configurations, the strain and strain energy density fields are essentially 0 within the stiff fibers, and also show good convergence properties with no obvious spurious oscillations near the transition from matrix to fiber. We notice that in the case where the fibers are further apart, the energy density and strain fields suggest that the fibers can be essentially considered dilute and non-interacting. However, when the fibers are closer together – as is typical of experimental specimens [HSP⁺19, RRWN17] – there is a clear interaction between the fields, and the energy density and strain fields take their largest values in the region between the fibers.

5.5 RPM Applied to the Accelerated FFT Method

Two important advantages of the RPM approach are, first, that it uses the fixed-point iterations to identify the unstable subspace; and, second, once identified, it performs Newton iterations on the unstable subspace with fixed-point iterations on the *entire* space. This makes it easy to apply the RPM strategy to existing fixed-point methods with minimal changes to algorithm / code. We therefore exploit this to apply the RPM approach to the Accelerated FFT method which is also based on a fixed-point framework. Given that the Accelerated FFT method has much superior performance compared to the Classical FFT method, we might expect that a RPM-Accelerated FFT method likely performs even better. We examine this issue below.

In [MS14], it is shown that the Accelerated FFT method is a special case of the general polarization-based fixed-point method developed by [MB12]. Algorithm 3 is the general polarization-based method, where α and β can take any value within a range. For the choice $\alpha = \beta = 2$, the polarization-based method reduces to the Accelerated FFT method. In [MS14], they also discuss and compare the Accelerated

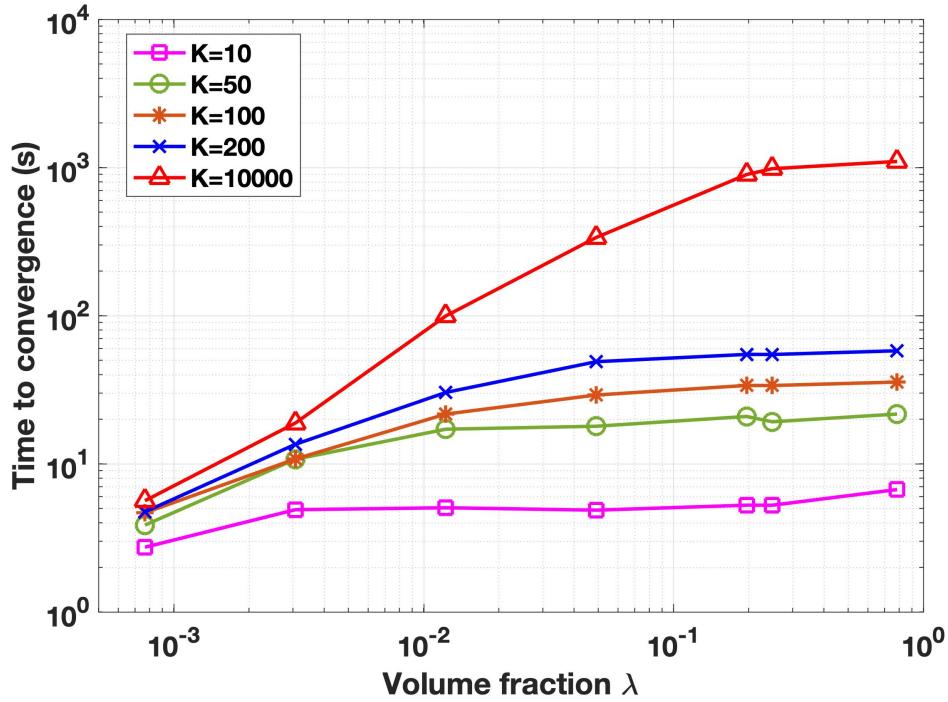


Figure 5: Time to convergence of RPM-FFT with varying λ and K . The lines connect the datapoints to aid visualization.

FFT method with other polarization-based methods, corresponding to the choices $\alpha = \beta = 1.5$ and $\alpha = \beta = 1$.

All of these methods are fixed-point based methods, and therefore require that the eigenvalues of the Jacobian are within the unit disk for convergence. The Accelerated FFT method has the lowest upper bound [MS14], and we therefore combine this with the RPM method.

In all of these methods, an important part of the method is the additive decomposition of the heterogeneous stiffness into a homogeneous reference medium and a fluctuation; see Section 4. The convergence rate is typically very sensitive to the choice of reference medium.

It has been observed that the choice of the reference medium affects the rate of convergence in many fixed-point based schemes. In the classical FFT method [MS98], the best choice was found to be the average of the supremum and infimum of the elastic moduli. In Accelerated FFT [MMS01] and Polarization FFT [MB12] the optimum is shown to be $E_o/E_f = \sqrt{K}$ when computing with a very high precision (10^{-10}).

However, as also discussed in Section 4, we expect the RPM-based approach to be insensitive to this choice. Therefore, we apply the RPM approach to the Accelerated FFT method, which is readily obtained by simply replacing the fixed-point operation in algorithm 2 with the iteration procedure in algorithm 3.

Table 1 compares the number of iterations and the time for Accelerated FFT and RPM-Accelerated FFT for a wide range of contrasts. The enhancement provided by RPM is increasingly large as the elastic contrast K increases.

We next examine the sensitivity to the choice of reference medium for Accelerated FFT and RPM-

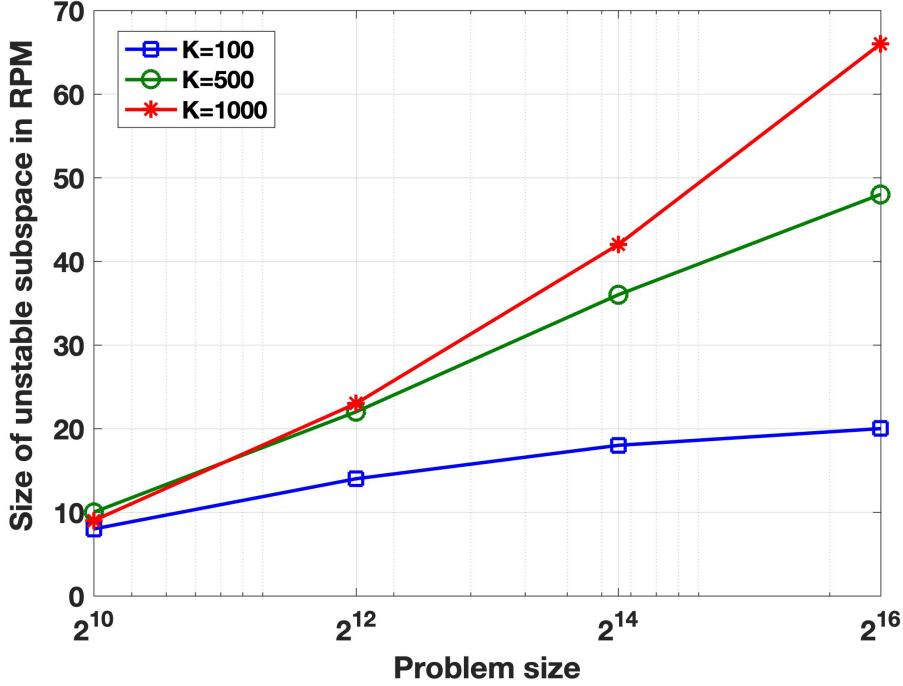


Figure 6: Size of unstable subspace with problem size. The lines connect the datapoints to aid visualization.

Algorithm 3 Polarization-FFT algorithm

- 1: **Iteration** $i + 1$ (given ϵ^i)
- 2: $\sigma^i \leftarrow \mathbf{C}(\mathbf{x}) : \epsilon^i$
- 3: $\mathbf{s}_a^i = \sigma^i + (1 - \beta)\mathbf{C}^0 : \epsilon^i$
- 4: $\mathbf{s}_b^i = \alpha\sigma^i + \beta\mathbf{C}^0 : \epsilon^i$
- 5: $\hat{\mathbf{s}}_b^i \leftarrow \mathcal{F}[\mathbf{s}_b^i];$
- 6: $\hat{\epsilon}_b^i \leftarrow -\hat{\mathbf{\Gamma}}^0 : \hat{\mathbf{s}}_b^i \quad \forall \mathbf{k} \neq \mathbf{0} \text{ and } \hat{\epsilon}_b(\mathbf{0})^i = \beta\mathbf{E}$
- 7: $\epsilon_b^i \leftarrow \mathcal{F}^{-1}[\hat{\epsilon}_b^i]$
- 8: $\epsilon^{i+1} = (\mathbf{C}(\mathbf{x}) + \mathbf{C}^0)^{-1} : (\mathbf{s}_a^i + \mathbf{C}^0 : \epsilon_b^i)$

Accelerated FFT. Fig. 8 compares the time to convergence and number of iterations – for a wide range of K – for the usual Accelerated FFT method against the RPM-Accelerated FFT method, for various choices of reference medium.

When K is small, the performance of the Accelerated FFT follows what has been observed in [MS14]: namely, the best choice is at when $E_o/E_f = \sqrt{K}$. But this is not the case for $K = 10^3, 10^4$. Meanwhile, RPM-Accelerated FFT is much less affected by the different choices of the reference over all of the K s being tested. As pointed out in [MS14], the $E_o/E_f = \sqrt{K}$ choice is valid when the precision is high (10^{-10}), which explains Fig. 8 (e,f,g,h), as all tests here are computed with 10^{-4} precision which is more in line with standard practice.

We notice, particularly at high contrast, that the use of RPM in conjunction with Accelerated FFT makes the method much more robust with regard to the choice of reference medium.

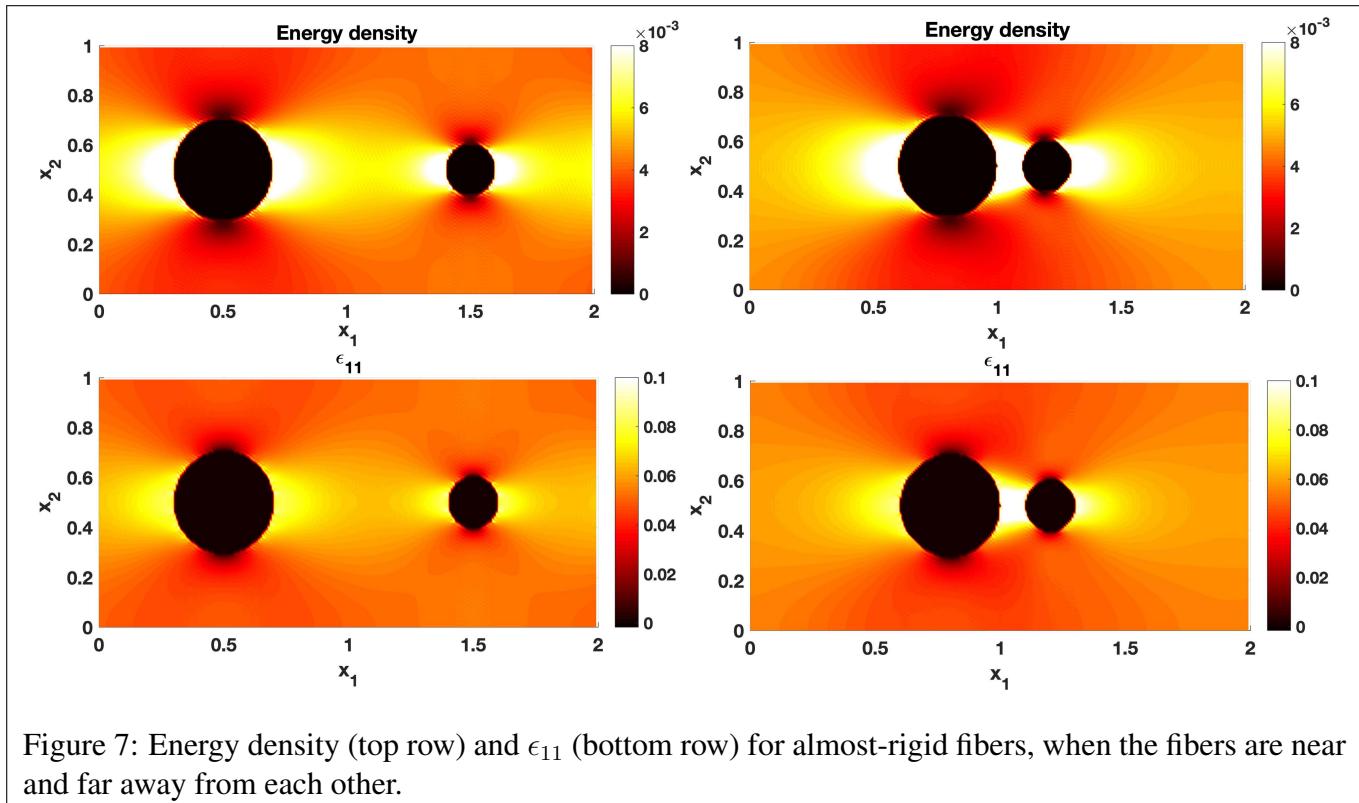


Figure 7: Energy density (top row) and ϵ_{11} (bottom row) for almost-rigid fibers, when the fibers are near and far away from each other.

6 Concluding Remarks

We have applied RPM [SK93] to the problem of determining the stress and strain in the unit cell of a periodic linear elastic material. Existing methods for this problem, e.g. Classical FFT [MS98] and Accelerated FFT [MMS01] methods, are fixed-point methods, and can have difficulty converging when the eigenvalues of the Jacobian lie outside the unit disk. While Newton methods can recover the convergence, they are very expensive on the high-dimensional problems of interest. RPM [SK93] provides an elegant and efficient balance between fixed-point and Newton methods: it uses the fixed-point iterations to adaptively identify the unstable subspace that requires Newton, and performs Newton only on that subspace. Fixed-point iterations can be performed on the complementary stable subspace. For practical reasons of not changing existing algorithms and code, fixed-point is performed on the entire space, but is

K	Number of Iterations		Time (seconds)	
	Acc-FFT	RPM-Acc-FFT	Acc-FFT	RPM-Acc-FFT
500	234	83	6.46	4.47
1000	335	81	8.51	4.42
5000	335	81	20.24	4.40
10000	1136	82	28.29	4.44
50000	2741	82	68.56	4.45
100000	3863	161	98.60	11.13

Table 1: Number of iterations and time to convergence for the Accelerated-FFT and RPM-Accelerated-FFT for a wide range of K .

then projected to the complementary stable subspace.

A variational perspective, using analogies from electromagnetism, provides insight into the reason that the RPM decomposition can work well. In particular, we expect that there are flat or unstable directions in the energy landscape in the context of certain formulations of heterogeneous linear elasticity, particularly when the stiffness tensor can vanish, as in the case of voids in materials. The RPM-FFT method proposed in this paper exploits this structure to perform Newton iterations in directions along which the fixed-point would be unstable.

Our results show that the RPM-FFT is more efficient than the fixed-point methods as the contrast is increased. Further, all of the fixed-point methods mentioned above require the use of a homogeneous reference medium as an important part of the method, but there is typically no systematic approach to selecting the reference medium. We see that while the fixed-point methods can be sensitive to the choice of reference medium, the RPM-FFT method is extremely robust in terms of convergence for a wide range of choices for the reference medium.

We also note the feature of RPM that it can be easily used to “wrap-around” any fixed-point algorithm. Consequently, while we have used it here for two specific instances of FFT-based methods, it can be easily combined with other formulations if they are found to be more efficient.

While we have applied RPM-FFT to the context of mechanical response, our overall approach is applicable to numerous linear problems with the same structure in homogenization [Mil02]. Further, while we have not examined nonlinear problems here, it is relatively straightforward to do this. Briefly, most methods for nonlinear problems using the FFT consist of the repeated use of linear solves [LR20], and the method presented can readily the other approaches in a modular way.

Finally, we note that while we have used Newton iterations for the space that is unstable under fixed-point iterations, it is possible to replace or combine Newton iterations with a conjugate gradient or other solver that could provide better performance [VZM14, GMC13, KBS14]. Potential improvement can be achieved by combining the model order reduction techniques for FFT solvers as proposed in [KMG⁺19] with RPM. This provides a direction for future exploration. In this context, we particularly highlight the recent interesting approach proposed in [Sch17], where a different variational principle is shown to provide promising results.

Acknowledgments

We thank Timothy Breitzman, Anthony Rollett, and Yi-Chung Shu for useful discussions. We acknowledge financial support from ARO Numerical Analysis (W911NF-17-1-0084), ONR Applied and Computational Analysis (N00014-18-1-2528), NSF DMREF Program (1628994), NSF Mechanics of Materials and Structures (1635407), and ARO MURI Program (W911NF-19-1-0245). We acknowledge the National Science Foundation for computing resources through the XSEDE program (TG-DMR120046) provided by Pittsburgh Supercomputing Center. Kaushik Dayal acknowledges an appointment to the National Energy Technology Laboratory Faculty Research Participation Program sponsored by the U.S. Department of Energy and administered by the Oak Ridge Institute for Science and Education.

Availability of Codes

The code developed and used for the calculations in this paper is available at

<https://github.com/KaushikDayalGroup/RPM-for-FFT-and-Elasticity>

References

- [AP13] M. Agoras and P. Ponte Castañeda. Iterated linear comparison bounds for viscoplastic porous materials with “ellipsoidal” microstructures. *Journal of the Mechanics and Physics of Solids*, 61(3):701–725, 2013.
- [BH16] IJ Beyerlein and A Hunter. Understanding dislocation mechanics at the mesoscale using phase field dislocation dynamics. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 374(2066):20150166, 2016.
- [Cia88] Philippe G Ciarlet. *Three-dimensional elasticity*, volume 20. Elsevier, 1988.
- [CS13] HZ Chen and YC Shu. Phase-field modeling of martensitic microstructure with inhomogeneous elasticity. *Journal of Applied Physics*, 113(12):123506, 2013.
- [EM99] DJ Eyre and GW Milton. A fast numerical scheme for computing the response of composites using grid refinement. *The European Physical Journal-Applied Physics*, 6(1):41–47, 1999.
- [Esh59] John Douglas Eshelby. The elastic field outside an ellipsoidal inclusion. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 252(1271):561–569, 1959.
- [GMC13] Lionel Gélébart and Romain Mondon-Cancel. Non-linear extension of fft-based methods accelerated by conjugate gradients to evaluate the mechanical behavior of composite materials. *Computational Materials Science*, 77:430–439, 2013.
- [GRL16] JT Graham, AD Rollett, and R LeSar. Fast fourier transform discrete dislocation dynamics. *Modelling and Simulation in Materials Science and Engineering*, 24(8):085005, 2016.
- [HSP⁺19] Spencer A. Hawkins, Ali Sarvestani, Xiaoyao Peng, Evan Sullivan, Evan Mungall, Kaushik Dayal, and Dhriti Nepal. In situ visualization of nanoscale damage and strain induced crystallization in polyurethane via atomic force microscopy. *submitted*, 2019.
- [JK90] Richard D James and David Kinderlehrer. Frustration in ferromagnetic materials. *Continuum Mechanics and Thermodynamics*, 2(3):215–239, 1990.
- [KBS14] Matthias Kabel, Thomas Böhlke, and Matti Schneider. Efficient fixed point and newton-kylov solvers for fft-based homogenization of elasticity at large deformations. *Computational Mechanics*, 54(6):1497–1514, 2014.
- [KMG⁺19] Julian Kochmann, Kiran Manjunatha, Christian Gierden, Stephan Wulffinghoff, Bob Svendsen, and Stefanie Reese. A simple and flexible model order reduction method for fft-based homogenization problems using a sparse sampling technique. *Computer Methods in Applied Mechanics and Engineering*, 347:622–638, 2019.

[Leb01] Ricardo A Lebensohn. N-site modeling of a 3d viscoplastic polycrystal using fast fourier transform. *Acta materialia*, 49(14):2723–2737, 2001.

[Lip93] Robert Lipton. Inequalities for electric and elastic polarization tensors with applications to random composites. *Journal of the Mechanics and Physics of Solids*, 41(5):809–833, 1993.

[LR20] Ricardo A Lebensohn and Anthony D Rollett. Spectral methods for full-field micromechanical modelling of polycrystalline materials. *Computational Materials Science*, 173:109336, 2020.

[MB12] Vincent Monchiet and Guy Bonnet. A polarization-based fft iterative scheme for computing the effective properties of elastic composites with arbitrary contrast. *International Journal for Numerical Methods in Engineering*, 89(11):1419–1436, 2012.

[Mil02] Graeme W Milton. The theory of composites. *The Theory of Composites*, by Graeme W. Milton, pp. 748. ISBN 0521781256. Cambridge, UK: Cambridge University Press, May 2002., page 748, 2002.

[MMS99] JC Michel, H Moulinec, and P Suquet. Effective properties of composite materials with periodic microstructure: a computational approach. *Computer methods in applied mechanics and engineering*, 172(1-4):109–143, 1999.

[MMS01] JC Michel, H Moulinec, and P Suquet. A computational scheme for linear and non-linear composites with arbitrary phase contrast. *International Journal for Numerical Methods in Engineering*, 52(1-2):139–160, 2001.

[MS94] Hervé Moulinec and Pierre Suquet. A fast numerical method for computing the linear and nonlinear mechanical properties of composites. *Comptes rendus de l'Académie des sciences. Série II, Mécanique, physique, chimie, astronomie*, 318(11):1417–1423, 1994.

[MS98] H Moulinec and P Suquet. A numerical method for computing the overall response of non-linear composites with complex microstructure. *Computer methods in applied mechanics and engineering*, 157(1-2):69–94, 1998.

[MS14] Hervé Moulinec and Fabrice Silva. Comparison of three accelerated fft-based schemes for computing the mechanical response of composite materials. *International Journal for Numerical Methods in Engineering*, 97(13):960–985, 2014.

[PL13] Marcel Porta and Turab Lookman. Heterogeneity and phase transformation in materials: Energy minimization, iterative methods and geometric nonlinearity. *Acta Materialia*, 61(14):5311–5340, 2013.

[PMB⁺20] Xiaoyao Peng, Nithin Mathew, Irene J Beyerlein, Kaushik Dayal, and Abigail Hunter. A 3d phase field dislocation dynamics model for body-centered cubic crystals. *Computational Materials Science*, 171:109217, 2020.

[PMG⁺15] S Piazolo, M Montagnat, F Grennerat, H Moulinec, and J Wheeler. Effect of local stress heterogeneities on dislocation fields: Examples from transient creep in polycrystalline ice. *Acta Materialia*, 90:303–309, 2015.

[RRWN17] Samit Roy, John Ryan, Samantha Webster, and Dhriti Nepal. A review of in situ mechanical characterization of polymer nanocomposites: Prospect and challenges. *Applied Mechanics Reviews*, 69(5):050802, 2017.

[SB01] YC Shu and K Bhattacharya. Domain patterns and macroscopic behaviour of ferroelectric materials. *Philosophical Magazine B*, 81(12):2021–2054, 2001.

[Sch17] Matti Schneider. An fft-based fast gradient method for elastic and inelastic unit cell homogenization problems. *Computer Methods in Applied Mechanics and Engineering*, 315:846–866, 2017.

[SK93] Gautam M Shroff and Herbert B Keller. Stabilization of unstable procedures: the recursive projection method. *SIAM Journal on numerical analysis*, 30(4):1099–1120, 1993.

[VZM14] Jaroslav Vondřejc, Jan Zeman, and Ivo Marek. An fft-based galerkin method for homogenization of periodic media. *Computers & Mathematics with Applications*, 68(3):156–173, 2014.

[YD11] Lun Yang and Kaushik Dayal. A completely iterative method for the infinite domain electrostatic problem with nonlinear dielectric media. *Journal of Computational Physics*, 230(21):7821–7829, 2011.

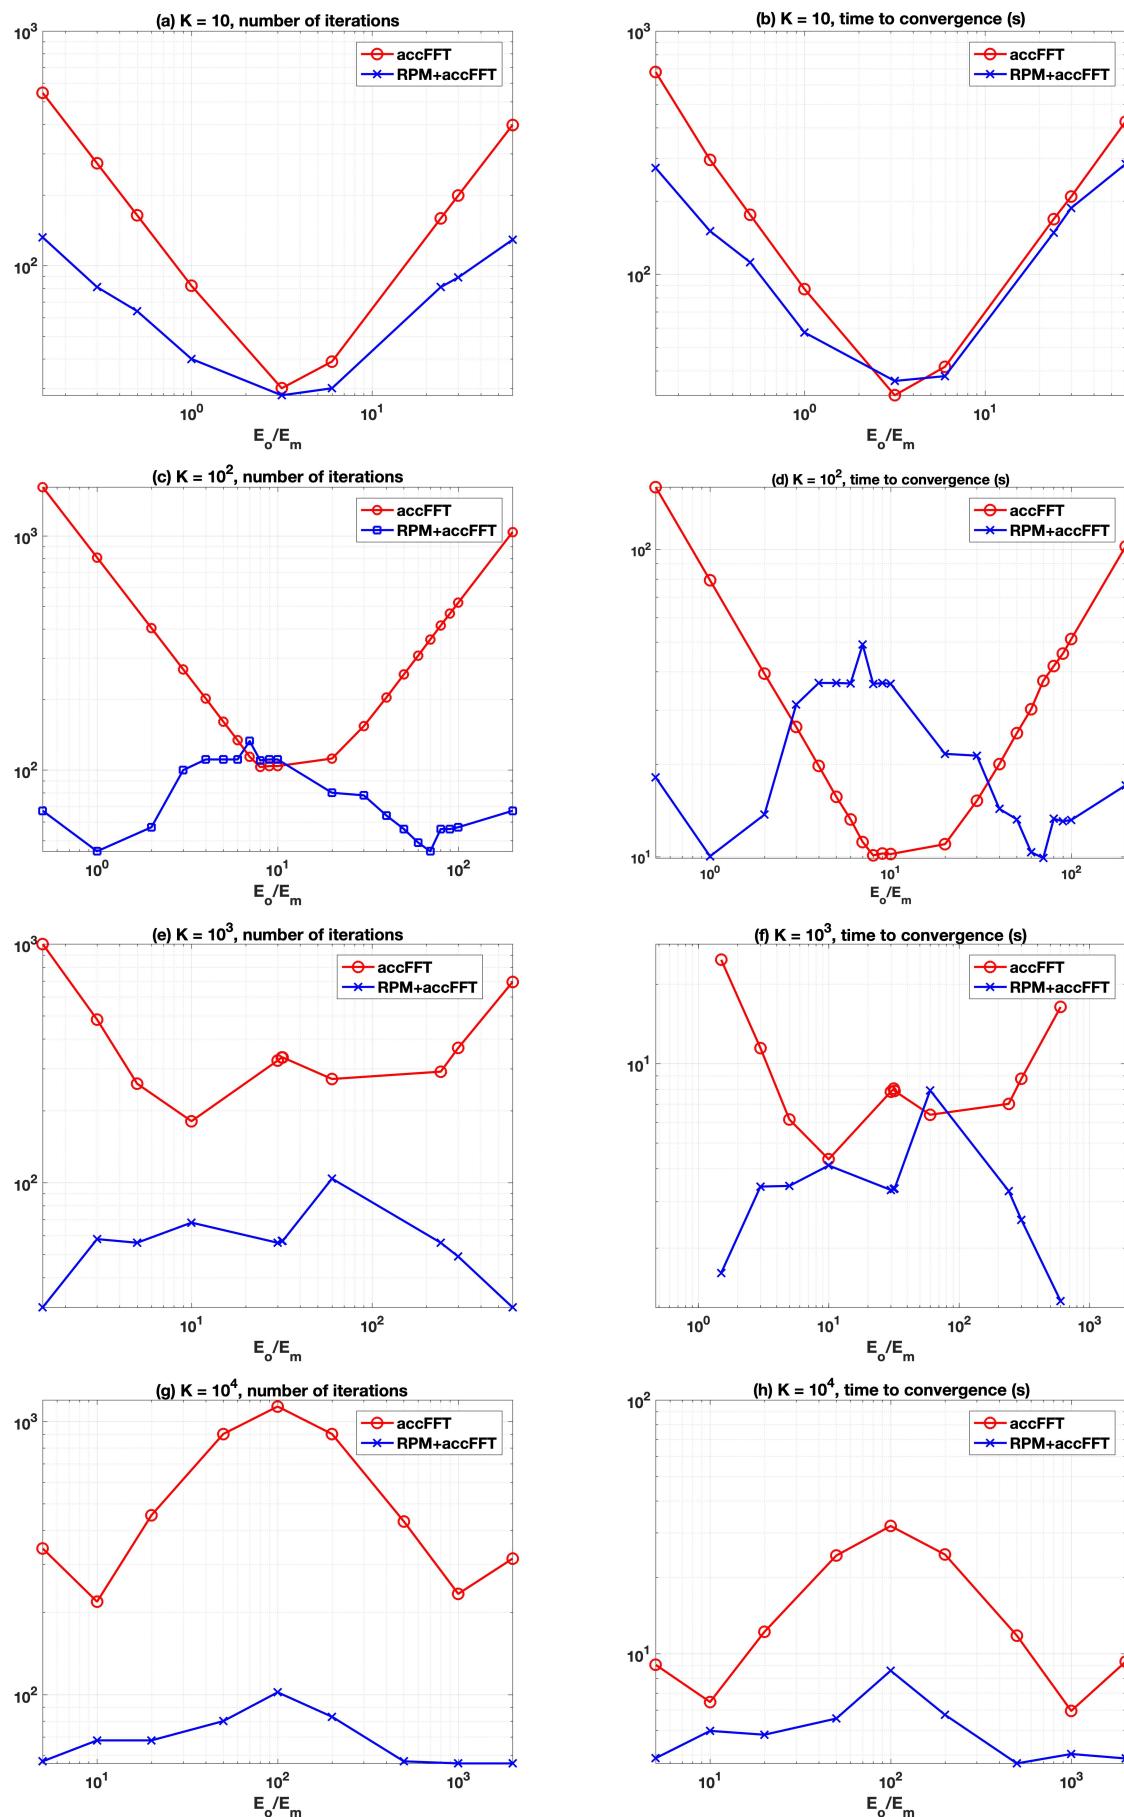


Figure 8: Number of iterations and times to convergence as a function of reference medium modulus E_o . We examine $K = 10, 10^2, 10^3, 10^4$ in the first, second, third and fourth rows respectively.