PhD Naval Architecture, Ocean and Marine Engineering Computational Homogenisation of Cracked Composite Materials using Ordinary State-Based Peridynamics

Wenxuan Xia

Peridynamics Research Center
Department of Naval Architecture, Ocean and Marine Engineering
University of Strathclyde, Glasgow

December 18, 2025

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

The copyright of this thesis belongs to the author under the terms of the United Kingdom Copyright Acts as qualified by University of Strathclyde Regulation 3.50. Due acknowledgement must always be made of the use of any material contained in, or derived from, this thesis.

Signed:

Date:

Abstract

This thesis presents a scalable computational homogenisation framework based on ordinary state-based peridynamics for the effective characterisation of cracked composite materials with periodic microstructures. Conventional continuum methods face challenges in modelling discontinuities such as matrix cracking and fibre-matrix debonding, whereas peridynamics naturally accommodates such damage through its integral formulation. However, integration with computational homogenisation in the ordinary state-based form remains limited.

To bridge this gap, a first-order homogenisation scheme is developed within a parallel `C++/MPI` environment. A key novelty is a raytracing-based bond-breaking algorithm that converts arbitrary cracks into triangle-mesh surfaces, enabling accurate interaction removal and damage tracking. Volumetric periodic boundary conditions are formulated for peridynamic RVEs to ensure displacement periodicity and force anti-periodicity, consistent with classical homogenisation. The framework supports fully 3D scalable simulations and employs harmonic mean sampling for interfacial bond properties to enhance physical fidelity.

Numerical studies on fibre-reinforced composite RVEs demonstrate accurate prediction of effective stiffness and local stress fields, validated against analytical and micromechanical models. The results show robust capture of complex crack morphologies and interactions while maintaining parallel efficiency. The proposed framework thus provides a general-purpose and scalable tool for peridynamic homogenisation of damaged composites, forming a foundation for future multiscale and inelastic material modelling.

Acknowledgements

This research was carried out under the guidance of my supervisors, Prof. Erkan Oterkus and Prof. Selda Oterkus. Their academic insight, encouragement, and continuous support were instrumental throughout this project.

I am deeply grateful to my parents for their unwavering support, patience, and belief in me. Their presence has been a constant source of strength and motivation.

I would also like to thank the Department of Naval Architecture, Ocean and Marine Engineering for supporting this research and covering the necessary high-performance computing resources.

Computational work was performed using the ARCHIE-WeSt High Performance Computing Centre, which provided the infrastructure essential for the simulations conducted in this study.

Contents

Contents

Contents

Contents

List of Figures

List of Figures

List of Figures

x

List of Tables

List of Tables

Chapter 1

Introduction

The effective behaviour of composite materials, particularly those with periodic microstructures, plays a critical role in the design and analysis of modern engineering systems. Accurate prediction of their macroscopic response requires computational models that capture heterogeneity, anisotropy, and damage mechanisms such as matrix cracking, fibre-matrix debonding, and crack interactions. Conventional continuum mechanics approaches, while well established, require special enrichment or remeshing techniques to represent such discontinuities. Extended finite element methods (XFEM), cohesive elements, and remeshing schemes have achieved notable success in this regard, yet they often involve complex algorithmic handling and mesh dependence when simulating cracks.

Peridynamics [1], a nonlocal reformulation of continuum mechanics expressed through integral equations, provides a more natural treatment of discontinuities. Fracture and debonding emerge directly from the governing equations without additional criteria, allowing seamless transition from intact to damaged states. While bond-based peridynamics has been extensively applied to fracture studies, its restrictive constitutive form, particularly the limitation on Poisson's ratio, reduces its suitability for general material modelling. The more general *ordinary state-based peridynamic* (OSB-PD) formulation [2] removes these restrictions, enabling anisotropy and complex stress states, and is therefore better suited for multiscale and homogenisation studies.

Despite its promise, the use of OSB-PD in computational homogenisation remains

limited. Previous studies have largely relied on simplified bond-based models, often restricted to two-dimensional RVEs and lacking realistic damage representation. Consequently, there are few demonstrations of scalable three-dimensional peridynamic homogenisation or quantitative performance benchmarks for large RVEs under crack-induced damage. Addressing these limitations requires both algorithmic innovation and computational efficiency.

This thesis presents a scalable computational homogenisation framework based on ordinary state-based peridynamics, explicitly designed to capture crack-induced damage within periodic composite microstructures. The work addresses two major challenges in the field. First, to overcome the modelling difficulty of complex discontinuities, a *raytracing-based bond-breaking algorithm* is introduced. This method converts arbitrary crack surfaces into triangle-mesh geometries for robust and geometrically consistent detection of broken bonds, enabling stable and physically meaningful tracking of interacting cracks and debonding interfaces. Second, to alleviate the heavy computational cost of 3D homogenisation, the framework is implemented in parallel using `C++` and MPI, supporting scalable domain decomposition and efficient large-scale simulation.

In addition, *volumetric periodic boundary conditions* [3, 4] are formulated for the peridynamic setting to ensure displacement periodicity and force anti-periodicity across opposing RVE faces. This ensures consistency with classical homogenisation theory while preserving the nonlocal character of peridynamics. The framework also employs harmonic mean sampling for interfacial bond properties to enhance the physical accuracy of composite modelling.

A series of numerical benchmarks on fibre-reinforced composite RVEs are conducted to validate the proposed method. Results demonstrate accurate prediction of effective stiffness and local stress distributions, consistent with analytical and micromechanical references. The framework efficiently captures complex crack morphologies and interactions while maintaining high scalability on parallel computing systems.

This document is organised as follows:

- Chapter 2 introduces the fundamentals of peridynamic theory, including kine-

matics, state definitions, and governing equations.

- Chapter 3 presents the theoretical formulation of the computational homogenisation procedure, including volume averaging and boundary condition implementation.

- Chapter 4 details the numerical implementation, solver configuration, and MPI-based parallelisation strategies.

- Chapter 5 reports benchmark results in two-dimensional settings, focusing on convergence, accuracy, and damage simulation.

- Chapter 6 extends the study to three-dimensional RVEs, demonstrating the framework's scalability and predictive capability.

- Chapter 7 concludes the thesis with a summary of findings and future research directions.

Through the integration of state-based peridynamics, periodic boundary enforcement, raytracing-based crack modelling, and distributed parallel computation, this work bridges composite microstructural damage with effective macroscopic behaviour. It provides a unified and scalable tool for peridynamic computational homogenisation, forming a foundation for future multiscale modelling of inelastic and damaged composites.

Chapter 2

Peridynamic Theory

In this chapter, we summarize the essential formulations of peridynamics, focusing on the bond-based and ordinary state-based models. The mathematical definitions and material parameters for these models are discussed in detail, based primarily on the foundational works by Silling and collaborators [1, 2, 5, 6] (see also [7]).

2.1 Introduction

Peridynamics is a nonlocal theory of solid mechanics introduced by Silling [1]. Unlike classical continuum mechanics, which relies on spatial derivatives of the displacement field, peridynamics formulates the equations of motion using integral expressions over a finite neighbourhood. This fundamental shift allows the theory to naturally accommodate discontinuities such as crack initiation and propagation, without requiring additional fracture criteria or enrichment techniques.

The region of influence around a material point is defined by a finite-length parameter called the *horizon*. Within this horizon, material points interact through pairwise or state-based force functions. As the horizon length tends to zero, classical elasticity can be recovered as a limiting case [8]. Conversely, at nanoscale regimes, peridynamics can incorporate interatomic interactions such as van der Waals forces by tailoring the response functions appropriately [9]. These properties position peridynamics as a multiscale-capable framework that bridges continuum and atomistic modelling.

4

The bond-based peridynamic model (BBPD), introduced in [5], assumes that material behaviour arises from pairwise forces between points within the horizon. While simple and computationally efficient, this formulation imposes limitations on Poisson's ratio, specifically, it is restricted to 1/4 for 3D and 2D plane strain models, and 1/3 for 2D plane stress models.

To overcome these restrictions, a more general framework known as state-based peridynamics was later developed [2, 6]. This formulation introduces the concept of vector and scalar states, enabling forces to depend on the collective deformation of a material point's neighbourhood. In particular, the ordinary state-based peridynamic (OSBPD) model allows a broader range of constitutive behaviours and material responses, including those with arbitrary Poisson's ratios.

2.2 Kinematics and State Definitions

Consider a solid body occupying a region $\Omega(t) \subset \mathbb{R}^d$ at time t in a d-dimensional Euclidean space, where $d \in \{1, 2, 3\}$. Each point $\mathbf{x} \in \Omega(t)$ represents a peridynamic particle associated with a differential volume dV' and a mass density $\rho(\mathbf{x})$. The undeformed, stress-free reference configuration at $t = 0$ is denoted by $\Omega_0 \subset \mathbb{R}^d$.

For a particle $\mathbf{x} \in \Omega_0$, interaction occurs with all particles $\mathbf{x}'$ such that the reference distance $|\mathbf{x}' - \mathbf{x}|$ lies within a finite range δ, referred to as the *horizon*. The vector connecting a pair of such interacting particles is termed a *bond*, and the collection of all such $\mathbf{x}'$ within the horizon of $\mathbf{x}$ is known as its *family*. This interaction domain is defined as

$$H_\mathbf{x} = \left\{ \mathbf{x}' \in \Omega_0 : |\mathbf{x}' - \mathbf{x}| \leq \delta \right\}. \tag{2.1}$$

The horizon $H_\mathbf{x}$ is commonly taken to be circular in 2D and spherical in 3D, although other horizon shapes may be adopted depending on the formulation or application.

2.2.1 Deformation Kinematics

Let $\mathbf{x}, \mathbf{x}' \in \Omega_0$ denote two particles in the reference configuration. At a later time $t > 0$, the body undergoes a deformation, and each particle experiences a displacement: $\mathbf{u}(\mathbf{x})$

and $\mathbf{u}(\mathbf{x}')$, respectively. The updated positions in the deformed configuration are given by

$$\mathbf{y} = \mathbf{x} + \mathbf{u}(\mathbf{x}), \tag{2.2a}$$

$$\mathbf{y}' = \mathbf{x}' + \mathbf{u}(\mathbf{x}'). \tag{2.2b}$$

The relative position vector in the deformed state becomes

$$\mathbf{y}' - \mathbf{y} = (\mathbf{x}' - \mathbf{x}) + \left[\mathbf{u}(\mathbf{x}') - \mathbf{u}(\mathbf{x})\right]. \tag{2.3}$$

The peridynamic *bond stretch*, a scalar quantity characterising elongation, is defined as

$$S = \frac{|\mathbf{y}' - \mathbf{y}| - |\mathbf{x}' - \mathbf{x}|}{|\mathbf{x}' - \mathbf{x}|}. \tag{2.4}$$

2.2.2 State-Based Representation

To facilitate generalisation and improve mathematical clarity, the concept of *state-based* peridynamics was introduced in [2]. In this formulation, physical quantities are represented by operators known as *vector states*, which act on bonds to yield kinematic or force responses.

The **deformation state**, denoted $\underline{\mathbf{Y}}$, maps the reference bond vector to the deformed bond vector as follows

$$\mathbf{y}' - \mathbf{y} = \underline{\mathbf{Y}}(\mathbf{x}, t)\langle \mathbf{x}' - \mathbf{x}\rangle. \tag{2.5}$$

Similarly, the **force state**, $\underline{\mathbf{T}}$, returns the force density vector $\mathbf{t}$ exerted by a particle at $\mathbf{x}'$ on a particle at $\mathbf{x}$ as

$$\mathbf{t}(\mathbf{y}' - \mathbf{y}, \mathbf{x}' - \mathbf{x}, t) = \underline{\mathbf{T}}(\mathbf{x}, t)\langle \mathbf{x}' - \mathbf{x}\rangle. \tag{2.6}$$

These operator-based definitions facilitate the modelling of complex material behaviour, including anisotropy and microstructural heterogeneity.

2.3 Governing Equations and Constitutive Models

2.3.1 Equation of Motion

The peridynamic equation of motion, following the formulations by $[1, 2, 5]$, can be expressed in the following two equivalent forms.

$$\rho(\mathbf{x})\ddot{\mathbf{u}}(\mathbf{x}, t) = \int_{H_{\mathbf{x}}} \left(\mathbf{t}(\mathbf{y}' - \mathbf{y}, \mathbf{x}' - \mathbf{x}, t) - \mathbf{t}'(\mathbf{y} - \mathbf{y}', \mathbf{x} - \mathbf{x}', t) \right) dV' + \mathbf{b}(\mathbf{x}, t) \qquad (2.7a)$$

$$\rho(\mathbf{x})\ddot{\mathbf{u}}(\mathbf{x}, t) = \int_{H_{\mathbf{x}}} \left(\underline{\mathbf{T}}(\mathbf{x}, t)\langle \mathbf{x}' - \mathbf{x}\rangle - \underline{\mathbf{T}}(\mathbf{x}', t)\langle \mathbf{x} - \mathbf{x}'\rangle \right) dV' + \mathbf{b}(\mathbf{x}, t) \qquad (2.7b)$$

Here, $\mathbf{x}$ and $\mathbf{x}'$ denote the positions of material points in the reference configuration, while $\mathbf{y}$ and $\mathbf{y}'$ are their respective positions in the deformed configuration. The variable t represents time. $\rho(\mathbf{x})$ denotes the mass density, $\ddot{\mathbf{u}}(\mathbf{x}, t)$ the acceleration, dV' the differential volume element, and $\mathbf{b}(\mathbf{x}, t)$ the body force density. The peridynamic force densities $\mathbf{t}$ and $\mathbf{t}'$ have physical units of F/L^{2d}, where F and L are the units of force and length, respectively.

For static analysis, typical in the context of homogenisation, the governing equation simplifies to an equilibrium condition

$$\int_{H_{\mathbf{x}}} \left(\mathbf{t}(\mathbf{y}' - \mathbf{y}, \mathbf{x}' - \mathbf{x}, t) - \mathbf{t}'(\mathbf{y} - \mathbf{y}', \mathbf{x} - \mathbf{x}', t) \right) dV' + \mathbf{b}(\mathbf{x}, t) = \mathbf{0}. \qquad (2.8)$$

This condition states that the total internal and external force acting on a particle at position $\mathbf{x}$ must sum to zero.

2.3.2 Bond-Based Formulation

In the bond-based peridynamic model, the force density vectors $\mathbf{t}$ and $\mathbf{t}'$ are assumed to be equal in magnitude and opposite in direction. This assumption inherently satisfies the balance of angular momentum. As illustrated in Figure 2.1, the force density can be written as follows

$$\mathbf{t}(\mathbf{y}' - \mathbf{y}, \mathbf{x}' - \mathbf{x}, t) = \underline{\mathbf{T}}(\mathbf{x}, t)\langle \mathbf{x}' - \mathbf{x}\rangle = \frac{1}{2}C\frac{\mathbf{y}' - \mathbf{y}}{|\mathbf{y}' - \mathbf{y}|} = \frac{1}{2}\mathbf{f}(\mathbf{y}' - \mathbf{y}, \mathbf{x}' - \mathbf{x}, t), \qquad (2.9a)$$

Figure 2.1: Bond-based peridynamic force state. The force density vectors are shown in purple and are opposite in direction and equal in magnitude. The horizons in the undeformed and deformed configurations are indicated by red dashed lines.

$$\mathbf{t}'(\mathbf{y} - \mathbf{y}', \mathbf{x} - \mathbf{x}', t) = \underline{\mathbf{T}}(\mathbf{x}', t)\langle \mathbf{x} - \mathbf{x}'\rangle = -\frac{1}{2}C\frac{\mathbf{y}' - \mathbf{y}}{|\mathbf{y}' - \mathbf{y}|} = -\frac{1}{2}\mathbf{f}(\mathbf{y}' - \mathbf{y}, \mathbf{x}' - \mathbf{x}, t). \quad (2.9b)$$

Here, C is an auxiliary scalar parameter determined by material constants, the displacement field, and the peridynamic horizon. Given that $|\mathbf{t}| = |\mathbf{t}'|$, the equation of motion (2.7a) reduces to

$$\rho(\mathbf{x})\ddot{\mathbf{u}}(\mathbf{x}, t) = \int_{H_\mathbf{x}} \mathbf{f}(\mathbf{y}' - \mathbf{y}, \mathbf{x}' - \mathbf{x}, t)\,dV' + \mathbf{b}(\mathbf{x}, t). \qquad (2.10)$$

2.3.3   State-Based Formulation

In the ordinary state-based peridynamic model, the force density vectors $\mathbf{t}$ and $\mathbf{t}'$ are assumed to be opposite in direction but not necessarily equal in magnitude, as illustrated in Figure 2.2. This assumption also satisfies the balance of angular momentum. The force densities are expressed as follows

$$\mathbf{t}(\mathbf{y}' - \mathbf{y}, \mathbf{x}' - \mathbf{x}, t) = \underline{\mathbf{T}}(\mathbf{x}, t)\langle \mathbf{x}' - \mathbf{x}\rangle = \frac{1}{2}A\frac{\mathbf{y}' - \mathbf{y}}{|\mathbf{y}' - \mathbf{y}|}, \qquad (2.11a)$$

Figure 2.2: Ordinary state-based peridynamic force state. The force density vectors are shown in purple and are opposite in direction but not equal in magnitude. The horizons in the undeformed and deformed configurations are indicated by red dashed lines.

$$\mathbf{t}'(\mathbf{y} - \mathbf{y}', \mathbf{x} - \mathbf{x}', t) = \underline{\mathbf{T}}(\mathbf{x}', t)\langle \mathbf{x} - \mathbf{x}' \rangle = -\frac{1}{2} B \frac{\mathbf{y}' - \mathbf{y}}{|\mathbf{y}' - \mathbf{y}|}. \tag{2.11b}$$

As in the bond-based case, the auxiliary parameters A and B are determined by the material properties, displacement field, and horizon size.

2.4 Material Parameter Derivation

2.4.1 Strain Energy and Force Densities

The pairwise auxiliary parameters A and B in force density vectors $\mathbf{t}$ and $\mathbf{t}'$ can be expressed in terms of the strain energy densities W and W' as follows

$$A = \frac{2}{V'} \frac{\partial W}{\partial (|\mathbf{y}' - \mathbf{y}|)}, \tag{2.12a}$$

$$B = -\frac{2}{V} \frac{\partial W'}{\partial (|\mathbf{y} - \mathbf{y}'|)}. \tag{2.12b}$$

In this work, we consider only isotropic elastic materials, which leads to the following strain energy density function at particle position $\mathbf{x}$:

$$W = a\theta^2 + b \int_{H_\mathbf{x}} \omega \left(|\mathbf{y}' - \mathbf{y}| - |\mathbf{x}' - \mathbf{x}| \right)^2 dV'. \tag{2.13}$$

Here, the dilatation θ is given explicitly by

$$\theta = d \int_{H_\mathbf{x}} \omega \left(|\mathbf{y}' - \mathbf{y}| - |\mathbf{x}' - \mathbf{x}| \right) \left(\frac{\mathbf{y}' - \mathbf{y}}{|\mathbf{y}' - \mathbf{y}|} \cdot \frac{\mathbf{x}' - \mathbf{x}}{|\mathbf{x}' - \mathbf{x}|} \right) dV'. \tag{2.14}$$

The constants a, b, and d are peridynamic material parameters derived from conventional elastic constants, namely Young's modulus E and Poisson's ratio ν. The scalar-valued influence function ω scales the interaction between material points $\mathbf{x}$ and $\mathbf{x}'$ as a function of their reference separation $|\mathbf{x}' - \mathbf{x}|$ and the horizon size δ. As detailed in [7], the explicit form of ω for state-based peridynamics is given as

$$\omega = \frac{\delta}{|\mathbf{x}' - \mathbf{x}|}. \tag{2.15}$$

In this study, the peridynamic horizon δ is assumed to be proportional to the particle spacing Δx, typically $\delta = 3\Delta x$. This choice ensures sufficient overlap between neighbouring particles for stable force transmission while maintaining computational efficiency. As $\Delta x \to 0$ with $\delta/\Delta x$ held constant, the nonlocal formulation converges to the corresponding local elasticity model. Conversely, larger values of δ lead to smoother but effectively stiffer responses due to increased nonlocal interactions. The selected ratio $\delta/\Delta x = 3$ provides a good balance between accuracy and computational cost, as demonstrated in prior studies (e.g., [1]; [7]).

By substituting (2.14) into (2.13) and evaluating the derivatives in (2.12), we obtain the explicit expression for the state-based peridynamic force density as follows

$$\mathbf{t}(\mathbf{y}' - \mathbf{y}, \mathbf{x}' - \mathbf{x}, t) = 2\omega \left[d \left(\frac{\mathbf{y}' - \mathbf{y}}{|\mathbf{y}' - \mathbf{y}|} \cdot \frac{\mathbf{x}' - \mathbf{x}}{|\mathbf{x}' - \mathbf{x}|} \right) a\theta \right.$$
$$\left. + b \left(|\mathbf{y}' - \mathbf{y}| - |\mathbf{x}' - \mathbf{x}| \right) \right] \frac{\mathbf{y}' - \mathbf{y}}{|\mathbf{y}' - \mathbf{y}|}, \tag{2.16a}$$

$$\mathbf{t}'(\mathbf{y} - \mathbf{y}', \mathbf{x} - \mathbf{x}', t) = -2\omega' \left[d \left(\frac{\mathbf{y} - \mathbf{y}'}{|\mathbf{y} - \mathbf{y}'|} \cdot \frac{\mathbf{x} - \mathbf{x}'}{|\mathbf{x} - \mathbf{x}'|} \right) a\theta' \right.$$
$$\left. + b \left(|\mathbf{y} - \mathbf{y}'| - |\mathbf{x} - \mathbf{x}'| \right) \right] \frac{\mathbf{y}' - \mathbf{y}}{|\mathbf{y}' - \mathbf{y}|}. \tag{2.16b}$$

The auxiliary parameters A and B from (2.11) can thus be written as

$$A = 4\omega \left[d \left(\frac{\mathbf{y}' - \mathbf{y}}{|\mathbf{y}' - \mathbf{y}|} \cdot \frac{\mathbf{x}' - \mathbf{x}}{|\mathbf{x}' - \mathbf{x}|} \right) a\theta + b \left(|\mathbf{y}' - \mathbf{y}| - |\mathbf{x}' - \mathbf{x}| \right) \right], \tag{2.17a}$$

$$B = 4\omega' \left[d \left(\frac{\mathbf{y} - \mathbf{y}'}{|\mathbf{y} - \mathbf{y}'|} \cdot \frac{\mathbf{x} - \mathbf{x}'}{|\mathbf{x} - \mathbf{x}'|} \right) a\theta' + b \left(|\mathbf{y} - \mathbf{y}'| - |\mathbf{x} - \mathbf{x}'| \right) \right]. \tag{2.17b}$$

In bond-based peridynamics, consistency requires that $A = B$. This condition eliminates the dilatation terms θ and θ', which implies:

$$ad = 0. \tag{2.18}$$

The explicit bond-based peridynamic force density function from (2.9) can be written as

$$\mathbf{t}(\mathbf{y}' - \mathbf{y}, \mathbf{x}' - \mathbf{x}, t) = -\mathbf{t}'(\mathbf{y} - \mathbf{y}', \mathbf{x} - \mathbf{x}', t) = \frac{1}{2} \mathbf{f}(\mathbf{y}' - \mathbf{y}, \mathbf{x}' - \mathbf{x}, t)$$
$$= 2\omega b \left(|\mathbf{y}' - \mathbf{y}| - |\mathbf{x}' - \mathbf{x}| \right) \frac{\mathbf{y}' - \mathbf{y}}{|\mathbf{y}' - \mathbf{y}|}. \tag{2.19}$$

The corresponding auxiliary parameter C is

$$C = 4\omega b \left(|\mathbf{y}' - \mathbf{y}| - |\mathbf{x}' - \mathbf{x}| \right). \tag{2.20}$$

To complete the formulation, we derive explicit expressions for the parameters a, b, and d in terms of elastic constants by comparing the peridynamic strain energy density in (2.13) with the classical continuum mechanics counterpart.

2.4.2 Parameter Calibration for Three-Dimensional Isotropic Materials

The constitutive law for a linear isotropic elastic material is given by $\sigma_{ij} = C_{ijkl}\varepsilon_{kl}$, or in Voigt notation:

$$
\begin{bmatrix} \sigma_{xx} \\ \sigma_{yy} \\ \sigma_{zz} \\ \sigma_{yz} \\ \sigma_{zx} \\ \sigma_{xy} \end{bmatrix} = \begin{bmatrix} \lambda + 2\mu & \lambda & \lambda & 0 & 0 & 0 \\ \lambda & \lambda + 2\mu & \lambda & 0 & 0 & 0 \\ \lambda & \lambda & \lambda + 2\mu & 0 & 0 & 0 \\ 0 & 0 & 0 & \mu & 0 & 0 \\ 0 & 0 & 0 & 0 & \mu & 0 \\ 0 & 0 & 0 & 0 & 0 & \mu \end{bmatrix} \begin{bmatrix} \varepsilon_{xx} \\ \varepsilon_{yy} \\ \varepsilon_{zz} \\ 2\varepsilon_{yz} \\ 2\varepsilon_{zx} \\ 2\varepsilon_{xy} \end{bmatrix} . \tag{2.21}
$$

The Lamé parameters λ and μ relate to Young's modulus E and Poisson's ratio ν via

$$
\lambda = \frac{E\nu}{(1+\nu)(1-2\nu)}, \qquad \mu = \frac{E}{2(1+\nu)}. \tag{2.22}
$$

Simple Shear Analysis

To determine the peridynamic parameter b, consider a simple shear strain state

$$
\varepsilon = \begin{bmatrix} 0 & 0 & 0 & s & 0 & 0 \end{bmatrix}^T, \tag{2.23}
$$

this corresponds to the deformation gradient

$$
\boldsymbol{F} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & s \\ 0 & 0 & 1 \end{bmatrix} . \tag{2.24}
$$

The resulting stress vector is

$$
\boldsymbol{\sigma} = \mathbf{D}\varepsilon = \begin{bmatrix} 0 & 0 & 0 & \mu s & 0 & 0 \end{bmatrix}^T . \tag{2.25}
$$

The strain energy density from classical continuum mechanics is

$$
W_{shear}^{ccm} = \frac{1}{2}\boldsymbol{\sigma}^T\varepsilon = \frac{\mu s^2}{2} = \frac{E s^2}{4(1+\nu)}. \tag{2.26}
$$

Since simple shear produces no volumetric change, dilatation is

$$
\theta_{shear}^{ccm} = 0. \tag{2.27}
$$

Figure 2.3: Representation of the bond $\langle \mathbf{x}' - \mathbf{x} \rangle$ in the spherical coordinate system, with $\mathbf{x}$ as the origin.

We now evaluate the same deformation under the peridynamic formulation using spherical coordinates (r_s, θ_s, ϕ_s) as shown in Figure 2.3, with

$$\left| \mathbf{x}' - \mathbf{x} \right| = r_s, \tag{2.28a}$$

$$\mathbf{x}' - \mathbf{x} = \begin{bmatrix} r_s \cos \theta_s \sin \phi_s \\ r_s \sin \theta_s \sin \phi_s \\ r_s \cos \phi_s \end{bmatrix}. \tag{2.28b}$$

The deformed bond under simple shear becomes

$$\mathbf{y}' - \mathbf{y} = \boldsymbol{F} \cdot (\mathbf{x}' - \mathbf{x}) = \begin{bmatrix} r_s \cos \theta_s \sin \phi_s \\ r_s \sin \theta_s \sin \phi_s + s r_s \cos \phi_s \\ r_s \cos \phi_s \end{bmatrix}, \tag{2.29a}$$

$$\left| \mathbf{y}' - \mathbf{y} \right| \approx \left(1 + \frac{s \sin(2\phi_s) \sin(\theta_s)}{2} \right) r_s. \tag{2.29b}$$

The bond length is linearised for small shear strains $s \ll 1$ using a first-order Taylor expansion, with higher-order terms $O(s^2)$ neglected. The volume element in spherical

13

Chapter 2. Peridynamic Theory

coordinates is

$$dV' = r_s^2 \sin \phi_s \, d\phi_s \, d\theta_s \, dr_s. \tag{2.30}$$

The corresponding peridynamic strain energy density in (2.13) for simple shear is

$$
W_{shear}^{pd} = b \int_{H_\mathbf{x}} \omega \left(|\mathbf{y}' - \mathbf{y}| - |\mathbf{x}' - \mathbf{x}| \right)^2 dV'
$$

$$
= b \int_0^\delta \int_0^{2\pi} \int_0^\pi \frac{\delta}{r_s} \left\{ \left[1 + \frac{s \sin(2\phi_s) \sin(\theta_s)}{2} \right] r_s - r_s \right\}^2 r_s^2 \sin \phi_s \, d\phi_s \, d\theta_s \, dr_s. \tag{2.31}
$$

Exploiting isotropy of the spherical horizon, all odd angular terms vanish upon integration, and the angular integrals reduce to a constant factor. The strain energy therefore becomes

$$
W_{shear}^{pd} = b \frac{4\pi}{15} s^2 \int_0^\delta \delta \, r_s^3 \, dr_s, \tag{2.32}
$$

integrating yields:

$$
W_{shear}^{pd} = \frac{\pi \delta^5 s^2}{15} b. \tag{2.33}
$$

Equating (2.33) with the classical strain energy density (2.26) gives

$$
b = \frac{15\mu}{2\pi\delta^5} = \frac{15E}{4\pi\delta^5(1+\nu)}. \tag{2.34}
$$

Isotropic Expansion Analysis

To derive parameters a and d, consider the isotropic expansion

$$
\varepsilon = \begin{bmatrix} s & s & s & 0 & 0 & 0 \end{bmatrix}^T, \tag{2.35}
$$

this corresponds to the deformation gradient

$$
\mathbf{F} = \begin{bmatrix} 1+s & 0 & 0 \\ 0 & 1+s & 0 \\ 0 & 0 & 1+s \end{bmatrix}. \tag{2.36}
$$

14

Chapter 2. Peridynamic Theory

The resulting stress vector is

$$\boldsymbol{\sigma} = \begin{bmatrix} (3\lambda + 2\mu)s \\ (3\lambda + 2\mu)s \\ (3\lambda + 2\mu)s \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} \frac{Es}{1-2\nu} \\ \frac{Es}{1-2\nu} \\ \frac{Es}{1-2\nu} \\ 0 \\ 0 \\ 0 \end{bmatrix}. \tag{2.37}$$

The classical strain energy density and dilatation become

$$W_{normal}^{ccm} = \frac{3(3\lambda + 2\mu)s^2}{2} = \frac{3Es^2}{2(1-2\nu)}, \tag{2.38}$$

$$\theta_{normal}^{ccm} = 3s. \tag{2.39}$$

The deformed bond under isotropic expansion becomes

$$\mathbf{y}' - \mathbf{y} = \boldsymbol{F} \cdot (\mathbf{x}' - \mathbf{x}) = \begin{bmatrix} r_s \cos\theta_s \sin\phi_s (1+s) \\ r_s \sin\theta_s \sin\phi_s (1+s) \\ r_s \cos\phi_s (1+s) \end{bmatrix}, \tag{2.40a}$$

$$|\mathbf{y}' - \mathbf{y}| = (1+s)r_s. \tag{2.40b}$$

The volume element in spherical coordinates is

$$dV' = r_s^2 \sin\phi_s \, d\phi_s \, d\theta_s \, dr_s. \tag{2.41}$$

The corresponding peridynamic strain energy density in (2.13) for isotropic expansion is

$$W_{normal}^{pd} = a(\theta_{normal}^{ccm})^2 + b \int_{H_\mathbf{x}} \omega \left(|\mathbf{y}' - \mathbf{y}| - |\mathbf{x}' - \mathbf{x}| \right)^2 dV'$$
$$= a(3s)^2 + b \int_0^\delta \int_0^{2\pi} \int_0^\pi \frac{\delta}{r_s} \left[(1+s)\, r_s - r_s \right]^2 r_s^2 \sin\phi_s \, d\phi_s \, d\theta_s \, dr_s. \tag{2.42}$$

Since the integrand is independent of angular variables, the integration over the spher-

ical horizon reduces to a radial integral, with the angular contribution yielding a factor of 4π. Integrating yields

$$W^{pd}_{normal} = 9as^2 + \pi\delta^5 bs^2. \tag{2.43}$$

The corresponding peridynamic dilatation in (2.14) for isotropic expansion is

$$
\begin{aligned}
\theta^{pd}_{normal} &= d \int_{H_\mathbf{x}} \omega\left(|\mathbf{y}' - \mathbf{y}| - |\mathbf{x}' - \mathbf{x}|\right) \left(\frac{\mathbf{y}' - \mathbf{y}}{|\mathbf{y}' - \mathbf{y}|} \cdot \frac{\mathbf{x}' - \mathbf{x}}{|\mathbf{x}' - \mathbf{x}|}\right) dV' \\
&= d \int_0^\delta \int_0^{2\pi} \int_0^\pi \frac{\delta}{r_s} \left[(1 + s)\, r_s - r_s\right] \left(\frac{\boldsymbol{r_s}}{r_s} \cdot \frac{\boldsymbol{r_s}}{r_s}\right) r_s^2 \sin\phi_s \, d\phi_s \, d\theta_s \, dr_s.
\end{aligned}
\tag{2.44}
$$

Since $\mathbf{y}' - \mathbf{y}$ is collinear with $\mathbf{x}' - \mathbf{x}$ under isotropic expansion, the unit vectors are identical and $(\boldsymbol{r_s}/r_s) \cdot (\boldsymbol{r_s}/r_s) = 1$. The integrand is therefore independent of angular variables. Integrating yields

$$\theta^{pd}_{normal} = \frac{4\pi\delta^4 ds}{3}. \tag{2.45}$$

By equating (2.45) with (2.39), the expression for d is obtained as

$$d = \frac{9}{4\pi\delta^4}. \tag{2.46}$$

Similarly, equating (2.43) with (2.38) yields the relationship between a and b. Solving these equations gives

$$a = \frac{\lambda - \mu}{2} = \frac{E(4\nu - 1)}{4(1 + \nu)(1 - 2\nu)}. \tag{2.47}$$

Summary of Peridynamic Parameters in 3D

$$a = \frac{\lambda - \mu}{2} = \frac{E(4\nu - 1)}{4(1 + \nu)(1 - 2\nu)}, \tag{2.48a}$$

$$b = \frac{15\mu}{2\pi\delta^5} = \frac{15E}{4\pi\delta^5(1 + \nu)}, \tag{2.48b}$$

$$d = \frac{9}{4\pi\delta^4}. \tag{2.48c}$$

2.4.3 Parameter Calibration for Plane Stress

In two-dimensional plane stress conditions, the stress and strain vectors in Voigt notation reduce to the following expressions

$$\boldsymbol{\sigma} = \begin{bmatrix} \sigma_{xx} & \sigma_{yy} & \sigma_{xy} \end{bmatrix}^T, \tag{2.49a}$$

$$\boldsymbol{\varepsilon} = \begin{bmatrix} \varepsilon_{xx} & \varepsilon_{yy} & \varepsilon_{xy} \end{bmatrix}^T. \tag{2.49b}$$

The corresponding stiffness matrix, obtained by inverting the compliance matrix, is

$$\mathbf{D}^{ps} = \frac{E}{1-\nu^2} \begin{bmatrix} 1 & \nu & 0 \\ \nu & 1 & 0 \\ 0 & 0 & \frac{1-\nu}{2} \end{bmatrix}. \tag{2.50}$$

To determine the peridynamic material parameters under plane stress, we follow the same calibration procedure as in subsection 2.4.2 by considering two deformation modes: simple shear and isotropic expansion. In each case, we compute the classical strain energy density and match it to the corresponding peridynamic expression.

Simple Shear Analysis

To determine the peridynamic parameter b, consider a simple shear strain state

$$\boldsymbol{\varepsilon} = \begin{bmatrix} 0 & 0 & s \end{bmatrix}^T, \tag{2.51}$$

this corresponds to the deformation gradient

$$\boldsymbol{F} = \begin{bmatrix} 1 & s \\ 0 & 1 \end{bmatrix}. \tag{2.52}$$

The resulting stress vector is

$$\boldsymbol{\sigma} = \mathbf{D}^{ps}\boldsymbol{\varepsilon} = \begin{bmatrix} 0 & 0 & \frac{Es}{2(1+\nu)} \end{bmatrix}^T. \tag{2.53}$$

Chapter 2. Peridynamic Theory

The strain energy density from classical continuum mechanics is

$$W_{shear}^{ccm} = \frac{1}{2}\boldsymbol{\sigma}^T \boldsymbol{\varepsilon} = \frac{Es^2}{4(1+\nu)}. \tag{2.54}$$

Since simple shear produces no volumetric change

$$\theta_{shear}^{ccm} = 0. \tag{2.55}$$

We now evaluate the same deformation under the peridynamic formulation using polar coordinate (r_p, θ_p), with

$$\left|\mathbf{x}' - \mathbf{x}\right| = r_p, \tag{2.56a}$$

$$\mathbf{x}' - \mathbf{x} = \begin{bmatrix} r_p \cos\theta_p \\ r_p \sin\theta_p \end{bmatrix}. \tag{2.56b}$$

The deformed bond under simple shear becomes

$$\mathbf{y}' - \mathbf{y} = \boldsymbol{F} \cdot (\mathbf{x}' - \mathbf{x}) = \begin{bmatrix} r_p \cos\theta_p + sr_p \sin\theta_p \\ r_p \sin\theta_p \end{bmatrix}, \tag{2.57a}$$

$$\left|\mathbf{y}' - \mathbf{y}\right| \approx \left[1 + (\sin\theta_p \cos\theta_p)s\right] r_p. \tag{2.57b}$$

The bond length is linearised for small shear strains $s \ll 1$ using a first-order Taylor expansion, with higher-order terms $O(s^2)$ neglected. The volume element in polar coordinate (r_p, θ_p) is

$$dV' = r_p\, d\theta_p\, dr_p. \tag{2.58}$$

The corresponding peridynamic strain energy density in (2.13), expressed in integral form for simple shear with h denoting the thickness is

$$\begin{aligned} W_{shear}^{pd} &= b \int_{H_{\mathbf{x}}} \omega \left(\left|\mathbf{y}' - \mathbf{y}\right| - \left|\mathbf{x}' - \mathbf{x}\right|\right)^2 dV' \\ &= bh \int_0^\delta \int_0^{2\pi} \frac{\delta}{r_p} \left\{\left[1 + (\sin\theta_p \cos\theta_p)s\right] r_p - r_p\right\}^2 r_p\, d\theta_p\, dr_p. \end{aligned} \tag{2.59}$$

Exploiting isotropy of the circular horizon, the angular integration yields a constant factor, independent of radius. The strain energy therefore becomes

$$W_{shear}^{pd} = b\,\frac{\pi}{4}\,s^2\,h\int_0^\delta \delta\,r_s^2\,dr_s, \tag{2.60}$$

integrating yields

$$W_{shear}^{pd} = \frac{\pi h \delta^4 s^2}{12}b. \tag{2.61}$$

Equating (2.61) with the classical strain energy density (2.54) gives

$$b = \frac{3E}{\pi h \delta^4 (1+\nu)}. \tag{2.62}$$

Isotropic Expansion Analysis

To derive parameters a and d, consider the isotropic expansion

$$\varepsilon = \begin{bmatrix} s & s & 0 \end{bmatrix}^T, \tag{2.63}$$

this corresponds to the deformation gradient

$$\boldsymbol{F} = \begin{bmatrix} 1+s & 0 \\ 0 & 1+s \end{bmatrix}. \tag{2.64}$$

The resulting stress vector is

$$\boldsymbol{\sigma} = \mathbf{D}^{ps}\varepsilon = \begin{bmatrix} \frac{Es}{1-\nu} \\ \frac{Es}{1-\nu} \\ 0 \end{bmatrix}. \tag{2.65}$$

The classical strain energy density and dilatation become

$$W^{ccm}_{normal} = \frac{1}{2}\boldsymbol{\sigma}^T\boldsymbol{\varepsilon} = \frac{Es^2}{1-\nu}, \tag{2.66}$$

$$\theta^{ccm}_{normal} = 2s. \tag{2.67}$$

The deformed bond under isotropic expansion becomes

$$\mathbf{y}' - \mathbf{y} = \boldsymbol{F} \cdot (\mathbf{x}' - \mathbf{x}) = \begin{bmatrix} r_p \cos\theta_p(1+s) \\ r_p \sin\theta_p(1+s) \end{bmatrix}, \tag{2.68a}$$

$$\left|\mathbf{y}' - \mathbf{y}\right| = (1+s)r_p. \tag{2.68b}$$

The volume element in polar coordinate (r_p, θ_p) is

$$dV' = r_p\,d\theta_p\,dr_p. \tag{2.69}$$

The corresponding peridynamic strain energy density in (2.13), expressed in integral form for isotropic expansion with h denoting the thickness is

$$\begin{aligned}
W^{pd}_{normal} &= a(\theta^{ccm}_{normal})^2 + b\int_{H_\mathbf{x}} \omega\left(\left|\mathbf{y}' - \mathbf{y}\right| - \left|\mathbf{x}' - \mathbf{x}\right|\right)^2 dV' \\
&= a(2s)^2 + bh\int_0^\delta \int_0^{2\pi} \frac{\delta}{r_p}\left[(1+s)r_p - r_p\right]^2 r_p\,d\theta_p\,dr_p.
\end{aligned} \tag{2.70}$$

Since the integrand is independent of θ_p, the angular integration yields a factor of 2π. Integrating yields

$$W^{pd}_{normal} = 4as^2 + \frac{2}{3}\pi bh\delta^4 s^2. \tag{2.71}$$

The corresponding peridynamic dilatation in (2.14), expressed in integral form for isotropic expansion with h denoting the thickness is

$$\begin{aligned}
\theta^{pd}_{normal} &= d\int_{H_\mathbf{x}} \omega\left(\left|\mathbf{y}' - \mathbf{y}\right| - \left|\mathbf{x}' - \mathbf{x}\right|\right)\left(\frac{\mathbf{y}' - \mathbf{y}}{\left|\mathbf{y}' - \mathbf{y}\right|} \cdot \frac{\mathbf{x}' - \mathbf{x}}{\left|\mathbf{x}' - \mathbf{x}\right|}\right) dV' \\
&= dh\int_0^\delta \int_0^{2\pi} \frac{\delta}{r_p}\left[(1+s)r_p - r_p\right]\left(\frac{\boldsymbol{r_p}}{r_p} \cdot \frac{\boldsymbol{r_p}}{r_p}\right) r_p\,d\theta_p\,dr_p.
\end{aligned} \tag{2.72}$$

Chapter 2. Peridynamic Theory

Since $\mathbf{y}' - \mathbf{y}$ is collinear with $\mathbf{x}' - \mathbf{x}$ under isotropic expansion, the unit vectors are identical and $(\boldsymbol{r_p}/r_p) \cdot (\boldsymbol{r_p}/r_p) = 1$. The integrand is therefore independent of angular variables. Integrating yields

$$\theta_{normal}^{pd} = \pi dh\delta^3 s. \tag{2.73}$$

By equating (2.67) with (2.73), the expression for d is obtained as

$$d = \frac{2}{\pi h \delta^3}. \tag{2.74}$$

Similarly, equating (2.66) with (2.71) yields the relationship between a and b. Solving these equations gives

$$a = \frac{E(3\nu - 1)}{4(1 - \nu)(1 + \nu)}. \tag{2.75}$$

Summary of Parameters in Plane Stress

$$a = \frac{E(3\nu - 1)}{4(1 - \nu)(1 + \nu)}, \tag{2.76a}$$

$$b = \frac{3E}{\pi h \delta^4 (1 + \nu)}, \tag{2.76b}$$

$$d = \frac{2}{\pi h \delta^3}. \tag{2.76c}$$

2.4.4 Parameter Calibration for Plane Strain

In the plane strain condition, deformation in the out-of-plane direction is assumed negligible. The constitutive relation is thus simplified to

$$\boldsymbol{\sigma} = \mathbf{D}^{pe}\boldsymbol{\varepsilon}, \tag{2.77}$$

where $\mathbf{D}^{pe}$ is the reduced 3×3 stiffness matrix under the plane strain assumption, defined as

$$\mathbf{D}^{pe} = \frac{E}{(1+\nu)(1-2\nu)} \begin{bmatrix} 1-\nu & \nu & 0 \\ \nu & 1-\nu & 0 \\ 0 & 0 & \frac{1-2\nu}{2} \end{bmatrix}. \tag{2.78}$$

To calibrate the peridynamic parameters, we first examine the system under a simple shear condition. Under such loading, the stress tensor $\boldsymbol{\sigma}$ and the corresponding strain energy density W_{shear}^{ccm} remain identical to those in the plane stress case. Therefore, parameter b is the same as in plane stress case:

$$b = \frac{3E}{\pi\delta^4(1+\nu)}. \tag{2.79}$$

Next, we consider the isotropic expansion in the plane. Under such loading condition, the stress state becomes

$$\boldsymbol{\sigma} = \mathbf{D}^{pe}\boldsymbol{\varepsilon} = \begin{bmatrix} \frac{Es}{(1+\nu)(1-2\nu)} \\ \frac{Es}{(1+\nu)(1-2\nu)} \\ 0 \end{bmatrix}. \tag{2.80}$$

The corresponding strain energy density under classical continuum mechanics is

$$W_{normal}^{ccm} = \frac{1}{2}\boldsymbol{\sigma}^T\boldsymbol{\varepsilon} = \frac{Es^2}{(1+\nu)(1-2\nu)}. \tag{2.81}$$

The dilatation in this case is also the same as that under plane stress:

$$\theta_{normal}^{ccm} = 2s, \tag{2.82}$$

which leads to a similar expression for parameter d:

$$d = \frac{2}{\pi\delta^3}. \tag{2.83}$$

The only parameter that differs between plane strain and plane stress is parameter a. It can be determined by equating the classical strain energy density with its peridynamic

counterpart, using the expression for d in (2.83):

$$a = \frac{E(4\nu - 1)}{4(1 + \nu)(1 - 2\nu)}. \tag{2.84}$$

Summary of Parameters in Plane Strain

$$a = \frac{E(4\nu - 1)}{4(1 + \nu)(1 - 2\nu)}, \tag{2.85a}$$

$$b = \frac{3E}{\pi\delta^4(1 + \nu)}, \tag{2.85b}$$

$$d = \frac{2}{\pi\delta^3}. \tag{2.85c}$$

A detailed derivation of the parameters a, b, and d under plane strain conditions can be found in Appendix B.

2.5 Discretisation Concepts

2.5.1 Discretisation of Peridynamic Equations

As shown in Figure 2.4, the peridynamic equation of motion (2.7) can be discretised to enable numerical computation. The resulting discrete form reads

$$\rho_k \ddot{\mathbf{u}}_k = \sum_{j=1}^{N} \left[\mathbf{t}_{kj}(\mathbf{y}_j - \mathbf{y}_k, \mathbf{x}_j - \mathbf{x}_k, t) - \mathbf{t}_{jk}(\mathbf{y}_k - \mathbf{y}_j, \mathbf{x}_k - \mathbf{x}_j, t) \right] V_j + \mathbf{b}_k, \tag{2.86a}$$

or equivalently,

$$\rho_k \ddot{\mathbf{u}}_k = \sum_{j=1}^{N} \left[\underline{\mathbf{T}}(\mathbf{x}_k, t)\langle \mathbf{x}_j - \mathbf{x}_k \rangle - \underline{\mathbf{T}}(\mathbf{x}_j, t)\langle \mathbf{x}_k - \mathbf{x}_j \rangle \right] V_j + \mathbf{b}_k. \tag{2.86b}$$

Here, N denotes the number of particles located within the horizon of particle k, and V_j is the nodal volume associated with particle j. From this point forward, the symbols k and j are reserved exclusively for particle indices. Unless otherwise noted, subscripts k and j refer to particle-specific properties, while the notations kj and jk denote bond-

Figure 2.4: Displacements of particles and the resulting force density. The horizon is set to three times the particle spacing, and bonds are rendered as grey dotted lines. Horizons in the undeformed and deformed configurations are illustrated.

specific properties between particle k and particle j.

2.5.2   Interface Treatment

The peridynamic parameters a, b, and d, as defined in (2.48), (2.76), and (2.85), were derived under the assumption of isotropic and homogeneous material domains. Consequently, these formulations are directly applicable only when the interacting particles belong to the same material phase.

However, in multi-material systems where distinct isotropic materials coexist, bonds may span across material interfaces. An even more complex scenario arises in functionally graded materials (FGMs), where engineering constants such as Young's modulus E and Poisson's ratio ν may vary continuously along the bond.

To handle such cases, several strategies are available for assigning effective bond properties that reflect the heterogeneous material environment. The following subsec-

tions describe and compare four representative approaches.

Direct Averaging of Engineering Constants

The direct averaging approach, as employed by [10, 11] in their studies on ordinary state-based peridynamic modelling of functionally graded materials (FGMs), assumes a straightforward treatment of material property variation. In this method, selected engineering material constants at the two bond endpoints (particles k and j) are averaged as

$$P_{kj} = \frac{P_k + P_j}{2}. \tag{2.87}$$

Here, P denotes engineering material constants such as Young's modulus E and Poisson's ratio ν. The averaged pair (E_{kj}, ν_{kj}) is then used in the constitutive mapping equations (2.48), (2.76), and (2.85) to determine the effective peridynamic parameters (a, b, d) associated with bond kj.

This method is computationally efficient and sufficiently accurate for smoothly varying materials but may introduce artefacts when applied near sharp interfaces.

Weaker-Dominant Assignment of Peridynamic Parameters

In scenarios where strength reduction at material interfaces is physically expected but computational efficiency is still a priority, one may opt for a conservative assignment strategy where the weaker material's properties are used to define the bond. This approach was adopted in [12].

Here, "weaker" is typically defined in terms of Young's modulus E, rather than peridynamic constants themselves, since the latter are not necessarily linearly correlated with material strength. The assignment is formulated as

$$P_{kj} = P_k \cdot \mathbb{K}(E_k \geq E_j) + P_j \cdot \mathbb{K}(E_k < E_j), \tag{2.88}$$

where $\mathbb{K}(\cdot)$ is the indicator function, equal to 1 when the condition is satisfied and 0 otherwise. Once the selected property pair P_{kj} is determined, it is passed to (2.48), (2.76), and (2.85) to compute the bond's effective peridynamic parameters (a, b, d).

Figure 2.5: Peridynamic bond kj at a dissimilar material interface, divided evenly into M segments with sampling points marked.

This method reflects a physically motivated lower-bound estimate while avoiding the overhead of fine-grained sampling.

Harmonic Mean of Engineering Constants

In scenarios where computational cost is not a primary constraint, the material constants along a bond can be evaluated using a *harmonic mean* of the engineering constants (e.g., Young's modulus E, Poisson's ratio ν) sampled along the bond path. This approach improves accuracy when modelling material interfaces between dissimilar constituents.

As shown in Figure 2.5, a random bond kj is divided into a total of M segments, each with length d_m. At the midpoint of each segment is a sampling point where material attributes are evaluated. The symbol $c^{(k,j;m)}$ represents the selected engineering constant sampled at point m along bond kj.

The harmonic mean over the M sampled points is computed as

$$c_{kj} = \left(\sum_{m=1}^{M} \frac{d_m}{c^{(k,j;m)}} \bigg/ \sum_{m=1}^{M} d_m \right)^{-1}. \tag{2.89}$$

This effective constant c_{kj} is then used to determine the peridynamic bond parameters (a, b, d) via the constitutive relations in (2.48), (2.76), and (2.85).

While the harmonic-mean sampling strategy provides more accurate representation of gradual or abrupt transitions across material interfaces, it is computationally more expensive than simpler averaging schemes. In many practical cases, particularly for homogeneous or only mildly heterogeneous media, the improvement in accuracy is

Figure 2.6: Peridynamic fictitious boundary layer with a depth equal to the horizon size δ.

marginal and often does not justify the additional cost. Nevertheless, for consistency and to ensure reliable treatment of strongly heterogeneous interfaces, this harmonic-mean parameter assignment is used in all subsequent benchmark tests within the proposed peridynamic computational homogenisation framework, unless otherwise stated.

Harmonic Mean of Peridynamic Parameters

In contrast to averaging engineering constants, one may also apply a harmonic mean directly to the peridynamic parameters (a, b, d) along a bond, each parameter is treated independently. This approach uses the same formulation as in (2.89), but with $c^{(k,j;m)}$ representing the peridynamic parameters sampled at point m along the bond between particles k and j.

This direct treatment avoids the intermediate step of converting engineering constants into peridynamic coefficients and was used in the author's earlier work [13]. While simpler to implement in certain contexts, this method can be less physically interpretable when dealing with complex heterogeneous materials or interface regions.

2.5.3 Fictitious Boundary

Peridynamic equation of motion (2.7) does not contain any spatial derivatives, hence its explicit solution generally requires no constraint conditions. However, for those

situations where displacement boundary conditions are required, a fictitious boundary layer Ω^c as shown in Figure 2.6, could be defined along the boundary of domain Ω. According to [14], the depth of Ω^c should be equal to the horizon δ for an appropriate reflection of the boundary condition onto the problem domain.

2.5.4 Damage Modelling

Damage modelling in peridynamics is typically done by scaling the bond force density expression (2.16) with an activation function. Defined as a binary switch, it takes the form

$$\psi_{kj} = \begin{cases} 0, & \text{if bond intersects with discontinuities} \\ 1, & \text{otherwise} \end{cases} \tag{2.90}$$

Therefore, when a bond meets the criterion for breaking (i.e., when nonlocal damage is detected for that bond), its contribution to the force density in the equation of motion is set to zero.

Two-Dimensional Crack

In two-dimensional analyses, cracks and bonds are both represented as line segments. The bond breaking criteria in this case becomes the problem of determining whether two line segments intersect in a two-dimensional plane.

This is a fundamental problem in computational geometry. One possible approach to this problem is by utilising a vector based approach to find the intersection point, and determine if intersection happens within both line segments(and if not exist, parallel). Note that the intersection point's coordinate is redundant in this case as we are only interested in finding whether intersection happens. Therefore, a more efficient approach is to use cross product to check line orientations, as described in [15]

Consider bond AB and crack CD, where points A, B, C and D are endpoints of line segments with corresponding position vector $\mathbf{a}(a_x, a_y)$, $\mathbf{b}(b_x, b_y)$, $\mathbf{c}(c_x, c_y)$ and $\mathbf{d}(d_x, d_y)$, respectively. The intersection condition can be understood as points A and B on the opposite side of CD, as well as points C and D on the opposite side of AB. This can

be achieved by defining an orientation operator

$$\text{orient}(\mathbf{a}, \mathbf{b}, \mathbf{c}) = \text{cross}(\mathbf{b} - \mathbf{a}, \mathbf{c} - \mathbf{a}), \tag{2.91}$$

where $\text{cross}(\cdot)$ is the cross operator concerning only the magnitude and sign. For vectors $\boldsymbol{v}(v_x, v_y)$ and $\boldsymbol{w}(w_x, w_y)$, operator $\text{cross}(\cdot)$ is defined as

$$\text{cross}(\boldsymbol{v}, \boldsymbol{w}) = v_x w_y - v_y w_x. \tag{2.92}$$

Positive outcome of $\text{orient}(\mathbf{a}, \mathbf{b}, \mathbf{c})$ indicates path A to B to C forms a left turn, negative outcome indicates right turn, while zero means points A, B and C form a straight line.

Proper intersection condition of AB and CD can then be written as

$$\text{orient}(\mathbf{c}, \mathbf{d}, \mathbf{a}) \cdot \text{orient}(\mathbf{c}, \mathbf{d}, \mathbf{b}) < 0$$
$$\wedge \ \text{orient}(\mathbf{a}, \mathbf{b}, \mathbf{c}) \cdot \text{orient}(\mathbf{a}, \mathbf{b}, \mathbf{d}) < 0. \tag{2.93}$$

Note that condition (2.93) only considers proper intersection, while edge cases are disregarded.

Three-Dimensional Crack

In three dimensions, pre-existing cracks are represented as triangulated surface meshes. Each crack is defined by a location vector $\mathbf{x} = (x, y, z)$ and an orientation described by a quaternion $\mathbf{q} = (a, b, c, w)$. The crack surface is first defined in a local coordinate system, where each quadrilateral patch is decomposed into two triangles, $(c0, c1, c3)$ and $(c1, c2, c3)$. The global Cartesian coordinates of the mesh nodes are obtained by rotating the local coordinates using the quaternion $\mathbf{q}$ and then translating by $\mathbf{x}$.

Determining whether a peridynamic bond $(\mathbf{x}, \mathbf{x}')$ intersects a crack surface is therefore equivalent to checking for segment-triangle intersections. This work employs the Möller-Trumbore algorithm [16], which provides an efficient and widely validated method for ray-triangle intersection. For each triangle, the bond is treated as a finite-length ray, and an intersection is registered only if (i) the parametric intersection lo-

Figure 2.7: Illustration of randomly oriented three-dimensional crack surfaces generated within the RVE.

Figure 2.8: Example of a triangulated crack-surface mesh subjected to a prescribed quaternion-based rotation.

cation lies within the triangle, and (ii) the intersection lies between the two bond end-points. To avoid unnecessary computation, the algorithm is implemented in a clipped form in which intersection tests are terminated once the potential intersection falls outside the bond segment limits.

Compared with the two-dimensional case, several geometric edge cases arise in three dimensions, for example nearly coplanar configurations, grazing intersections, and intersection points lying exactly on triangle edges. These cases are handled explicitly by thresholding the determinant term of the Möller-Trumbore test and by disregarding ambiguous configurations that fall below numerical tolerances.

This ray-tracing approach naturally extends to crack surfaces of arbitrary complexity. Provided that the surface can be discretised into a triangle mesh using any suitable meshing algorithm, the same intersection procedure applies triangle-by-triangle. To assist reproducibility, a simplified pseudo-code version of the bond-crack intersection routine is included in Algorithm 1: This approach provides a general, robust, and

---

Algorithm 1 Bond-Crack Intersection Detection

1: **function** BONDINTERSECTSCRACK($bond, crack\_mesh$)
2:     **for** each triangle $tri \in crack\_mesh.triangles$ **do**
3:         $(hit, t) \leftarrow$ MollerTrumbore($bond.start, bond.length, bond.direction, tri$)
4:         **if** hit **and** $0 < t < 1$ **then**
5:             **return** $true$
6:     **return** $false$

---

highly efficient mechanism for bond breaking in three-dimensional peridynamic simulations involving arbitrarily oriented and arbitrarily shaped crack surfaces.

2.5.5 Surface Correction

Peridynamic parameters a, b and d are derived by calculating dilatation and strain energy density of a particle fully embedded within an isotropic homogeneous medium. Previously we considered the situation for dissimilar material interfaces, however, the requirement for complete horizon still stands. Thus, the previously derived peridynamic parameters are still not valid at free surfaces or material interfaces. To resolve this issue we resort to a surface correction procedure for parameter b and d. Note that parameter

a does not depend on horizon integration thus is exempt from surface correction.

The surface correction for parameter b and d covers the situation where particle is near free surfaces (including crack surfaces) or material interfaces. For computational homogenisation, the only free surfaces in the system are from cracks and other discontinuities, while the domain boundary is normally covered by periodic boundary condition which is considered full horizon computationally for particles reside within.

The surface correction factor for parameter b at particle k, denoted $s_{k,m}^{b}$ ($m = x, y, z$), can be obtained by comparing strain energy density definition under classical continuum mechanics against its peridynamic counterpart.

To eliminate dilatation from peridynamic strain energy density formulation (2.13), three simple shear deformation cases can be applied to the system separately for correction factors along each axial direction. They are written in terms of the deformation gradient as

$$\boldsymbol{F}_x = \begin{bmatrix} 1 & s & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \boldsymbol{F}_y = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & s \\ 0 & 0 & 1 \end{bmatrix}, \quad \boldsymbol{F}_z = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ s & 0 & 1 \end{bmatrix}, \tag{2.94}$$

with s being a small value. Deformation cases $\boldsymbol{F}_x$, $\boldsymbol{F}_y$ and $\boldsymbol{F}_z$ correspond to correction factor components $s_{k,x}^{b}$, $s_{k,y}^{b}$ and $s_{k,z}^{b}$, respectively. We then have

$$s_{k,m}^{b} = \frac{W_{k,m}^{ccm}}{W_{k,m}^{pd}} = \frac{\frac{1}{2}\mu s^2}{b \sum_{j=1}^{N} \omega_{kj}(|\mathbf{y}_j - \mathbf{y}_k| - |\mathbf{x}_j - \mathbf{x}_k|)^2 V_j}, \quad (m = x, y, z) \tag{2.95}$$

Similarly, surface correction factor for parameter d at particle k, denoted $s_{k,m}^{d}$ ($m = x, y, z$), can be obtained by applying three normal strain deformations separately for each axis. Written in deformation gradients

$$\boldsymbol{F}_x = \begin{bmatrix} 1+s & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \boldsymbol{F}_y = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1+s & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \boldsymbol{F}_z = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1+s \end{bmatrix}, \tag{2.96}$$

with s being a small value. Deformation cases $\boldsymbol{F}_x$, $\boldsymbol{F}_y$ and $\boldsymbol{F}_z$ correspond to correction

factor components $s_{k,x}^d$, $s_{k,y}^d$ and $s_{k,z}^d$, respectively. By comparing dilatation obtained using classical continuum mechanics against its peridynamic counterpart (2.14), we have

$$s_{k,m}^d = \frac{\theta_{k,m}^{ccm}}{\theta_{k,m}^{pd}} = \frac{s}{d \sum_{j=1}^{N} \omega_{kj}(|\mathbf{y}_j - \mathbf{y}_k| - |\mathbf{x}_j - \mathbf{x}_k|) \left(\frac{\mathbf{y}_j - \mathbf{y}_k}{|\mathbf{y}_j - \mathbf{y}_k|} \cdot \frac{\mathbf{x}_j - \mathbf{x}_k}{|\mathbf{x}_j - \mathbf{x}_k|} \right) V_j},$$

$$(m = x, y, z)$$

$$(2.97)$$

So far, we have obtained surface correction factor $s_{k,m}^b$ and $s_{k,m}^d$ for particle k. Peridynamic parameters b and d are bond related attributes, therefore we could calculate the direct average of $s_{k,m}^\alpha$ and $s_{j,m}^\alpha$, $(\alpha = b, d)$

$$\bar{s}_{kj,m}^\alpha = \frac{s_{k,m}^\alpha + s_{j,m}^\alpha}{2},$$

$$(2.98)$$

and then apply ellipsoid rule to $\bar{s}_{kj,m}^\alpha$ gives the final form of surface correction factors for parameter b and d as follows

$$\bar{S}_{kj}^\alpha = \left[\left(\frac{n_x}{\bar{s}_{kj,x}^\alpha} \right)^2 + \left(\frac{n_y}{\bar{s}_{kj,y}^\alpha} \right)^2 + \left(\frac{n_z}{\bar{s}_{kj,z}^\alpha} \right)^2 \right]^{-\frac{1}{2}},$$

$$(2.99)$$

where n_x, n_y and n_z are components of the unit vector:

$$\boldsymbol{n} = \frac{\mathbf{x}_j - \mathbf{x}_k}{|\mathbf{x}_j - \mathbf{x}_k|}.$$

$$(2.100)$$

Dilatation (2.14) and strain energy density (2.13) can then be surface corrected as follows

$$\theta_k = d \sum_{j=1}^{N} \bar{S}_{kj}^d \omega_{kj}(|\mathbf{y}_j - \mathbf{y}_k| - |\mathbf{x}_j - \mathbf{x}_k|) \left(\frac{\mathbf{y}_j - \mathbf{y}_k}{|\mathbf{y}_j - \mathbf{y}_k|} \cdot \frac{\mathbf{x}_j - \mathbf{x}_k}{|\mathbf{x}_j - \mathbf{x}_k|} \right) V_j,$$

$$(2.101)$$

$$W_k = a\theta_k^2 + b \sum_{j=1}^{N} \bar{S}_{kj}^b \omega_{kj}(|\mathbf{y}_j - \mathbf{y}_k| - |\mathbf{x}_j - \mathbf{x}_k|)^2 V_j.$$

$$(2.102)$$

Chapter 3

Computational Homogenisation Framework

This chapter presents the computational homogenisation framework developed to link microscale material behaviour with macroscale mechanical response. A first-order homogenisation approach is adopted, under the assumption of clear scale separation between the macro- and microscales. Theoretical foundations are established through averaging theorems and the Hill-Mandel energy consistency condition, culminating in the formal definition of effective material properties.

Analytical bounds and micromechanical models are reviewed to contextualise numerical results. Particular emphasis is placed on the enforcement of periodic boundary conditions within a peridynamic setting, which is critical for ensuring consistency across scales. The chapter concludes with a detailed description of the numerical procedure used to compute the effective stiffness tensor, forming the basis of the multiscale analysis employed throughout this work.

Portions of the content are adapted from work previously published by the author during the course of this PhD research [13, 17].

Chapter 3. Computational Homogenisation Framework

3.1 Introduction

Computational homogenisation is a multiscale modelling technique designed to bridge the gap between microscopic material behaviour and the effective properties used in macroscopic continuum models. Rather than relying solely on empirical or analytical estimates, computational homogenisation derives macroscopic stress-strain relationships from detailed simulations of a representative volume element (RVE) of the microstructure. The RVE is subjected to suitable boundary conditions, and its mechanical response is averaged to extract effective material properties. This methodology is particularly well suited for heterogeneous materials with periodic or statistically homogeneous microstructures, where an RVE can be meaningfully defined.

The foundations of computational homogenisation were established by classical works such as those of Hill [18,19], Nemat-Nasser [20], and Suquet [21], who formalised the mathematical principles connecting microscopic and macroscopic quantities via averaging theorems and energy consistency conditions. With the advent of the finite element method and modern computing power, these concepts evolved into practical multiscale strategies, particularly in the form of first-order computational homogenisation [22,23]. This approach is now widely applied to composites, polycrystals, porous media, and other complex materials. Typically, each integration point in a macroscopic finite element model is coupled to a microscale boundary value problem whose solution provides the effective constitutive response [24–26].

In parallel with computational approaches, a range of analytical micromechanical models, such as the Voigt [27] and Reuss [28] bounds, the Hashin–Shtrikman variational principles [29], and the Mori–Tanaka method [30,31] have been developed to estimate effective properties from known phase properties and geometries. A central theoretical foundation for many of these models is Eshelby's solution for an ellipsoidal inclusion in an infinite medium [32], which introduced the concept of the Eshelby tensor and enabled analytical treatment of inclusion problems in elasticity. Building on this, mean-field approaches such as the self-consistent scheme [33] and its generalised extensions [34] offer improved accuracy for polycrystalline and random composites by treating each

phase as an inclusion embedded in an effective medium. While less flexible than full-field numerical schemes, these analytical models remain useful for rapid estimation, parameter studies, and validation. Some of them, including modified Mori–Tanaka formulations for transversely isotropic media (e.g., the Tandon-Weng model [35]) and semi-empirical relations such as the Halpin-Tsai equations [36], will be reviewed in this chapter and later employed for comparison in Chapters 5 and 6.

Despite its broad applicability, classical computational homogenisation encounters challenges in situations involving discontinuities, damage, or other nonlocal effects, where standard continuum mechanics may fail to capture essential behaviour. In such cases, *peridynamics* [1], which is a nonlocal theory of continuum mechanics, offers a compelling alternative. Its integral-based formulation allows for the modelling of cracks and other discontinuities without the need for special enrichment techniques or discontinuity tracking. These features make peridynamics an attractive candidate for governing microscale behaviour within multiscale frameworks.

This chapter presents a computational homogenisation methodology in which state-based peridynamics [2] is employed to model the microscale RVE, while classical continuum mechanics governs the macroscopic scale. The overall strategy includes a review of theoretical foundations, averaging theorems, boundary conditions, and analytical micromechanical models, culminating in a robust procedure for computing effective stiffness tensors. This methodology forms the basis for the multiscale analysis presented in the subsequent chapters.

3.2   Multiscale Homogenisation Framework

The foundation of computational homogenisation lies in the assumption of a clear separation between the microscopic and macroscopic scales. This section formalises the scale separation hypothesis, outlines the adopted first-order homogenisation framework, and establishes the associated notational conventions.

3.2.1   Scale Separation and First-Order Homogenisation

A heterogeneous material that appears statistically homogeneous at the macroscale is considered. A well-defined scale separation is assumed between the microscopic and macroscopic length scales as follows

$$\ell_{\text{micro}} \ll \ell_{\text{macro}}.$$

This assumption ensures that macroscopic fields such as strain and stress vary smoothly over the domain. Consequently, a first-order computational homogenisation framework is adopted, wherein higher-order coupling effects and explicit microscale field gradients are neglected.

3.2.2   Multiscale Modelling Strategy

Figure 3.1: Computational homogenisation scheme

The computational homogenisation procedure adopts a two-scale modelling strategy. As illustrated in Figure 3.1, the macroscopic continuum domain is discretised using finite elements, with each integration point associated with a corresponding RVE at the microscale.

At each macroscopic integration point, prescribed boundary conditions, typically six independent small-strain components in the case of three-dimensional first-order homogenisation are imposed on the RVE. A boundary value problem (BVP) is solved over

the RVE, and the resulting homogenised response, such as effective stress or stiffness, is passed back to the macroscopic model to inform the constitutive behaviour.

3.2.3 Representative Volume Element (RVE)

The RVE must be sufficiently large to statistically capture the material's microstructural features, yet small enough relative to the macroscopic domain to ensure computational tractability. For materials with periodic microstructures, as considered in this work, the RVE can be reduced to a single periodic unit cell.

At the microscale, the displacement field $\boldsymbol{u}(\boldsymbol{x})$ is additively decomposed as

$$\boldsymbol{u}(\boldsymbol{x}) = \overline{\boldsymbol{\varepsilon}} \cdot \boldsymbol{x} + \tilde{\boldsymbol{u}}(\boldsymbol{x}), \tag{3.1}$$

where $\overline{\boldsymbol{\varepsilon}}$ is the prescribed macroscopic strain tensor, and $\tilde{\boldsymbol{u}}(\boldsymbol{x})$ denotes the displacement fluctuation field. For periodic media, $\tilde{\boldsymbol{u}}(\boldsymbol{x})$ is required to be periodic across opposing RVE boundaries. Rigid body motions are excluded, as the RVE displacement field is constructed to satisfy kinematic admissibility conditions, such as periodic boundary enforcement.

3.2.4 Notation Conventions

Throughout this chapter, the following notational conventions are used:

- The Einstein summation convention is adopted: repeated indices imply summation unless stated otherwise.

- Latin indices i, j, k, l denote Cartesian components and range from 1 to n, where n is the number of spatial dimensions (typically $n = 2$ or 3).

- Macroscopic quantities are denoted using an overline (e.g., $\overline{\boldsymbol{\varepsilon}}$, $\overline{\boldsymbol{\sigma}}$).

- Tensor and vector quantities are written in boldface, while scalar quantities are written in standard (non-bold) typeface.

- The double contraction between tensors is denoted by a colon symbol (:). For example, $\boldsymbol{A} : \boldsymbol{B} = A_{ijkl}B_{ijkl}$ for fourth-order tensors, or $\boldsymbol{\sigma} : \boldsymbol{\varepsilon} = \sigma_{ij}\varepsilon_{ij}$ for second-order tensors.

3.3 Volume Averages and Averaging Theorems

This section presents the mathematical foundations of volume averaging used in computational homogenisation, establishing the link between microscopic and macroscopic field quantities.

3.3.1 Volume Averages of Field Quantities

Volume averaging provides the mathematical foundation for relating microscopic fields to effective macroscopic quantities.

For a microscopic field quantity $f(\boldsymbol{x})$ defined over the RVE domain Ω_{RVE}, its volume average is defined as

$$\langle f \rangle = \frac{1}{|\Omega_{\mathrm{RVE}}|} \int_{\Omega_{\mathrm{RVE}}} f(\boldsymbol{x})\,\mathrm{d}V, \tag{3.2}$$

where $|\Omega_{\mathrm{RVE}}|$ denotes the volume (or area in 2D) of the RVE. These volume averages are used to define macroscopic field quantities. In particular, the macroscopic strain $\overline{\boldsymbol{\varepsilon}}$ and stress $\overline{\boldsymbol{\sigma}}$ are defined as the respective averages of their microscopic counterparts.

3.3.2 Averaging Theorems and Hill-Mandel Condition

Averaging theorems provide a rigorous framework to link microscale behaviour to macroscopic response. A central concept is the Hill-Mandel macrohomogeneity condition, which ensures energy consistency between the micro- and macro-scales.

Fundamental contributions by Hill [18], Nemat-Nasser and Hori [20], and Suquet [21] establish the conditions under which macroscopic fields can be obtained as volume averages of microscopic fields. These conditions require clear scale separation, a representative microstructure, and appropriate boundary conditions.

The essential averaging relations are:

- The macroscopic strain is the volume average of the microscopic strain field:

$$\overline{\varepsilon}_{ij} = \langle \varepsilon_{ij}. \rangle \tag{3.3}$$

- The macroscopic stress is the volume average of the microscopic stress field:

$$\overline{\sigma}_{ij} = \langle \sigma_{ij} \rangle. \tag{3.4}$$

- The Hill-Mandel condition ensures energetic equivalence between scales under virtual work:

$$\langle \sigma_{ij} \, \delta\varepsilon_{ij} \rangle = \overline{\sigma}_{ij} \, \delta\overline{\varepsilon}_{ij}. \tag{3.5}$$

Physically, the Hill-Mandel condition states that the work done by microscopic stresses on microscopic strains within a representative volume element (RVE) must equal the work done by the corresponding macroscopic stress on the macroscopic strain. In other words, the RVE behaves energetically like a homogeneous material at the macro-scale. This is particularly relevant for peridynamics, where nonlocal interactions replace classical stress-strain relations: enforcing the Hill-Mandel condition ensures that the computed nonlocal internal forces in the RVE consistently reproduce the macroscopic energy, allowing accurate multiscale predictions.

3.4 Analytical Estimates of Effective Properties

Before resorting to full computational homogenisation, several analytical approaches can provide estimates of the effective mechanical properties of heterogeneous materials. These methods, based on simplifying assumptions about the distribution of stress and strain within the microstructure, serve as useful benchmarks, offering bounds or approximate values that aid in preliminary design, parametric studies, and validation of numerical models.

Table 3.1 summarises the main analytical methods, which serve as benchmark references in later chapters, highlighting the type of estimate each one provides. The

following subsections describe each method in detail.

Table 3.1: Summary of analytical estimates of effective properties.

Method	Type of Estimate
Voigt and Reuss	Bounds
Hashin–Shtrikman	Bounds
Mori–Tanaka	Approximation
Tandon–Weng	Approximation
Halpin–Tsai	Empirical

3.4.1 Voigt and Reuss Bounds

A straightforward estimate of the effective stiffness tensor is given by the Voigt model [27], which assumes a uniform strain field throughout the RVE as:

$$\varepsilon_{ij}^{(r)} = \overline{\varepsilon}_{ij}, \tag{3.6}$$

where r indexes the material phases. Under this assumption, microscale displacement fluctuations are neglected, and equilibrium equations within the RVE need not be solved.

The effective stiffness tensor is obtained by a volume-weighted average of the individual phase stiffness tensors as:

$$\overline{C}_{ijkl}^{\text{Voigt}} = \sum_r f^{(r)} C_{ijkl}^{(r)}, \tag{3.7}$$

where $f^{(r)}$ is the volume fraction and $C_{ijkl}^{(r)}$ is the stiffness tensor of phase r.

The Voigt estimate provides an upper bound on the effective stiffness, assuming perfect strain compatibility and no local stress concentrations. It typically overestimates the true effective stiffness, especially when stiff phases dominate.

In contrast, the Reuss model [28] assumes a uniform stress field:

$$\sigma_{ij}^{(r)} = \overline{\sigma}_{ij}. \tag{3.8}$$

Here, stress continuity is enforced, but strain compatibility across phase interfaces is

not. Internal equilibrium equations are again not required.

The effective compliance tensor is computed as the volume-weighted average of the phase compliance tensors:

$$\overline{S}_{ijkl}^{\text{Reuss}} = \sum_r f^{(r)} S_{ijkl}^{(r)}, \tag{3.9}$$

where $S_{ijkl}^{(r)} = \left(C_{ijkl}^{(r)}\right)^{-1}$ is the compliance tensor of phase r. The effective stiffness tensor is then obtained:

$$\overline{C}_{ijkl}^{\text{Reuss}} = \left(\overline{S}_{ijkl}^{\text{Reuss}}\right)^{-1}. \tag{3.10}$$

The Reuss estimate yields a lower bound on effective stiffness, representing an idealised scenario of uniform stress and maximum strain heterogeneity. It often underestimates the overall stiffness, particularly in composites dominated by soft phases.

Taken together, the Voigt and Reuss bounds define an envelope within which the actual effective stiffness of the heterogeneous medium is expected to lie. More refined estimates (e.g., Hashin–Shtrikman) aim to narrow this interval.

3.4.2 Hashin–Shtrikman Bounds

The Hashin–Shtrikman (HS) bounds [29] provide rigorous theoretical estimates for the effective elastic moduli of two-phase composites with isotropic constituents. Derived using variational principles, these bounds are the tightest possible given only the volume fractions and elastic properties of the individual phases.

Consider a composite composed of two isotropic phases with bulk moduli K_1, K_2, shear moduli G_1, G_2, and volume fractions f_1 and $f_2 = 1 - f_1$. The Hashin–Shtrikman

bounds on the effective bulk modulus $\overline{K}$ and shear modulus $\overline{G}$ are given by:

$$\overline{K}^{\text{HS-}} = K_1 + \frac{f_2}{\dfrac{1}{K_2 - K_1} + \dfrac{3f_1}{3K_1 + 4G_1}}, \tag{3.11}$$

$$\overline{K}^{\text{HS+}} = K_2 + \frac{f_1}{\dfrac{1}{K_1 - K_2} + \dfrac{3f_2}{3K_2 + 4G_2}}, \tag{3.12}$$

$$\overline{G}^{\text{HS-}} = G_1 + \frac{f_2}{\dfrac{1}{G_2 - G_1} + \zeta_1 f_1}, \tag{3.13}$$

$$\overline{G}^{\text{HS+}} = G_2 + \frac{f_1}{\dfrac{1}{G_1 - G_2} + \zeta_2 f_2}, \tag{3.14}$$

where the auxiliary parameters ζ_i are defined as

$$\zeta_i = \frac{6(K_i + 2G_i)}{5G_i(3K_i + 4G_i)}, \quad i = 1, 2. \tag{3.15}$$

These bounds are tighter than the Voigt and Reuss estimates, which assume uniform strain and uniform stress, respectively. The Hashin–Shtrikman bounds are especially valuable when microstructural details are unavailable, but the constituent moduli and volume fractions are known.

In practice, these bounds serve as useful benchmarks for validating numerically computed effective properties. Deviations outside the HS bounds typically indicate numerical errors or a violation of the underlying assumptions of homogenisation theory.

3.4.3 Mori–Tanaka Method

The Mori–Tanaka method is a widely used mean-field homogenisation approach for estimating the effective elastic properties of composites with aligned inclusions embedded in a matrix [30, 31]. Unlike the Voigt, Reuss, and Hashin–Shtrikman bounds, which provide upper and lower limits, the Mori–Tanaka method yields a single estimate based on the assumption that the inclusions are sufficiently dispersed and interact with the average field of the matrix.

Consider a two-phase composite consisting of an isotropic matrix phase (denoted

by subscript 1) and ellipsoidal inclusions of an isotropic second phase (denoted by 2), with respective volume fractions f_1 and $f_2 = 1 - f_1$. The effective stiffness tensor $\overline{\boldsymbol{C}}$ is estimated using the Mori–Tanaka method as:

$$\overline{\boldsymbol{C}} = \boldsymbol{C}_1 + f_2 \left(\boldsymbol{C}_2 - \boldsymbol{C}_1\right) : \boldsymbol{A}, \tag{3.16}$$

where $\boldsymbol{A}$ is the strain concentration tensor, defined as:

$$\boldsymbol{A} = \left[\boldsymbol{I} + \boldsymbol{S} : \left(\boldsymbol{C}_1\right)^{-1} : \left(\boldsymbol{C}_2 - \boldsymbol{C}_1\right)\right]^{-1}. \tag{3.17}$$

Here, $\boldsymbol{I}$ is the fourth-order identity tensor, and $\boldsymbol{S}$ is the Eshelby tensor, also fourth-order, which depends on the shape of the inclusion and the elastic properties of the matrix phase.

Closed-form expressions for the effective bulk modulus $\overline{K}$ and shear modulus $\overline{G}$ can be derived from (3.16) in specific cases, depending on the inclusion aspect ratio and material contrast. Such formulations are detailed in [31]. These estimates typically lie within the Hashin–Shtrikman bounds.

The Mori–Tanaka method assumes that inclusions are non-interacting, uniformly dispersed, and embedded in an infinite matrix subjected to uniform remote loading. It performs well for composites with low to moderate inclusion volume fractions, especially when the inclusions are stiffer than the matrix. However, the method becomes less accurate as the inclusion content increases or when strong interactions between inclusions arise.

3.4.4 Tandon–Weng Model

The Tandon–Weng model [35] is a semi-analytical extension of the Mori–Tanaka method tailored for unidirectional composites reinforced by aligned, ellipsoidal inclusions. By incorporating Eshelby's tensorial solution for ellipsoidal inhomogeneities, the model explicitly accounts for inclusion shape, particularly the aspect ratio and its impact on anisotropic elastic behaviour. This approach makes the Tandon–Weng model especially suitable for fibre-reinforced composites, where the reinforcement phase can be idealised

as prolate spheroids aligned along a common axis.

The model considers a two-phase composite comprising an isotropic matrix (phase 1) and aligned ellipsoidal inclusions of a second isotropic phase (phase 2), with respective volume fractions f_1 and $f_2 = 1 - f_1$. Under the assumption of perfect bonding and linear elasticity, the effective stiffness tensor $\overline{C}$ exhibits transverse isotropy and admits closed-form expressions for all five independent engineering constants.

For instance, the effective longitudinal and transverse Young's moduli, $\overline{E}_1$ and $\overline{E}_2$, are expressed as:

$$\overline{E}_1 = \frac{E_1}{1 + f_2(A_1 + 2\nu_1 A_2)/A}, \tag{3.18}$$

$$\overline{E}_2 = \frac{E_1}{1 + f_2[-2\nu_1 A_3 + (1 - \nu_1)A_4 + (1 + \nu_1)A_5 A]/(2A)}, \tag{3.19}$$

where E_1 and ν_1 are the Young's modulus and Poisson's ratio of the matrix, respectively, and the geometric factors A and A_n are functions of the inclusion aspect ratio derived from the Eshelby tensor.

Analogous expressions are available for the in-plane and out-of-plane shear moduli $\overline{G}12$ and $\overline{G}23$, as well as the transverse bulk modulus $\overline{K}_{23}$, enabling a full characterisation of the transversely isotropic response. These closed-form relations underscore the sensitivity of transverse and shear properties to the inclusion geometry, even when the volume fraction remains fixed.

The Tandon–Weng model has found wide application in micromechanics and composite materials design, offering a convenient yet insightful tool for exploring how reinforcement shape and orientation influence effective stiffness. However, the model's applicability is limited by its assumptions, including uniform alignment, dilute concentrations, and non-interacting inclusions. It also assumes isotropic phases and perfect interfacial bonding, which may not hold in real systems with interfacial degradation or complex morphologies. Despite these simplifications, the model remains a foundational analytical method for estimating the geometry-dependent anisotropic properties of aligned composite systems.

3.4.5 Halpin–Tsai Equations

The Halpin–Tsai equations provide a semi-empirical framework for estimating the effective elastic properties of composite materials containing aligned, short or continuous fibres. Developed to bridge the gap between micromechanical theories and experimental results, the method introduces adjustable geometric parameters to account for reinforcement shape and orientation, offering engineers and materials scientists a flexible yet analytically tractable approach. Unlike purely theoretical models such as Mori–Tanaka or Eshelby-based solutions, the Halpin–Tsai equations are grounded in curve-fitting to experimental data and are particularly useful when dealing with composites reinforced by high-modulus fibres in a lower-modulus matrix.

The generalised Halpin–Tsai formulation is expressed as:

$$\overline{P} = P_m \cdot \frac{1 + \xi\eta f}{1 - \eta f}, \quad \text{with} \quad \eta = \frac{P_r/P_m - 1}{P_r/P_m + \xi}, \tag{3.20}$$

where $\overline{P}$ is the effective property of the composite (such as a modulus or shear property), P_r and P_m denote the reinforcement and matrix properties respectively, f is the volume fraction of the reinforcement, and ξ is a shape parameter that depends on the geometry and loading condition. The parameter ξ is often treated empirically; typical values range from 1 for disk-shaped fillers to 2 or more for aligned fibres.

The strength of the Halpin–Tsai model lies in its simplicity, adaptability, and ability to capture trends observed in experimental data across a variety of composite systems. It provides a convenient framework for quick estimates and parametric studies, especially in early-stage material design. However, its empirical nature also constitutes its primary limitation. The shape parameter ξ lacks a rigorous physical basis and must often be calibrated for specific systems. Additionally, the model assumes perfect bonding, isotropic constituents, and a uniform dispersion of inclusions, which may not hold in real composites. Despite these limitations, the Halpin–Tsai equations remain a widely used and practical tool for estimating effective composite properties in engineering applications.

3.5   Boundary Conditions at the Microscale

The choice of boundary conditions applied to the RVE plays a critical role in determining the computed effective properties. Different boundary conditions lead to distinct microscale responses and, consequently, different macroscopic estimates. The three most commonly employed types of boundary conditions in computational homogenisation are:

1. Prescribed displacement boundary conditions (Dirichlet),

2. Prescribed traction boundary conditions (Neumann),

3. Periodic boundary conditions (PBCs).

Each of these imposes different constraints on the RVE and corresponds to different assumptions regarding the surrounding material environment. Their implications are discussed in the subsections that follow.

3.5.1   Prescribed Displacements (Dirichlet)

Prescribed displacement boundary conditions, also known as Dirichlet boundary conditions, involve enforcing a macroscopic strain state by prescribing the displacement field along the boundary of the RVE. This approach requires solving the full microscale boundary value problem and, although computationally intensive, provides detailed resolution of the local mechanical response, including phase interactions and internal heterogeneity.

Formally, for a given macroscopic strain tensor $\bar{\boldsymbol{\varepsilon}}$, the displacement $\boldsymbol{u}(\boldsymbol{x})$ at any boundary point $\boldsymbol{x} \in \partial\Omega_{\text{RVE}}$ is prescribed as:

$$\boldsymbol{u}(\boldsymbol{x}) = \bar{\boldsymbol{\varepsilon}} \cdot \boldsymbol{x}, \quad \boldsymbol{x} \in \partial\Omega_{\text{RVE}}. \tag{3.21}$$

This condition enforces a uniform macroscopic strain across the RVE boundary while allowing the internal fields to develop freely in response to material heterogeneity. The resulting microscale stress field $\boldsymbol{\sigma}(\boldsymbol{x})$ is then volume-averaged to compute the

macroscopic stress $\overline{\boldsymbol{\sigma}}$. Due to the kinematic constraint imposed, this method yields an upper bound on the effective stiffness tensor, similar in nature to the Voigt estimate, though generally more accurate owing to the resolved microstructural detail.

3.5.2   Prescribed Tractions (Neumann)

In the Neumann boundary condition approach, tractions rather than displacements are prescribed along the boundary of the RVE. This corresponds to imposing a macroscopic stress state by applying consistent surface forces. For a given macroscopic stress tensor $\overline{\boldsymbol{\sigma}}$, the boundary traction $\boldsymbol{t}(\boldsymbol{x})$ at a point $\boldsymbol{x} \in \partial\Omega_{\text{RVE}}$ is specified as:

$$\boldsymbol{t}(\boldsymbol{x}) = \overline{\boldsymbol{\sigma}} \cdot \boldsymbol{n}(\boldsymbol{x}), \quad \boldsymbol{x} \in \partial\Omega_{\text{RVE}} \tag{3.22}$$

where $\boldsymbol{n}(\boldsymbol{x})$ denotes the outward unit normal vector at the boundary.

This condition ensures that the RVE is in equilibrium with the imposed macroscopic stress, but it does not constrain the displacement field, which remains undetermined up to a rigid body motion. Although the resulting microscale displacement field can still be used to compute an average strain, the absence of direct strain control makes this approach less compatible with first-order computational homogenisation schemes, where the macroscopic strain is typically the known input.

Due to this limitation, Neumann boundary conditions are generally unsuitable for multiscale frameworks that rely on strain-driven formulations. Accordingly, this boundary condition is not employed in the numerical implementation developed in this thesis.

3.5.3   Periodic Boundary Conditions for Peridynamics

While Dirichlet boundary conditions offer simple formulations for RVE problems, they may not fully capture the periodic nature of many heterogeneous materials. Periodic boundary conditions (PBCs) provide a more representative coupling between opposing boundaries of the RVE by enforcing continuity of displacement fluctuations and anti-periodicity of tractions. This approach not only ensures better consistency with the assumptions of first-order homogenisation but also minimises artificial boundary effects,

making it especially suitable for materials with repeating microstructures.
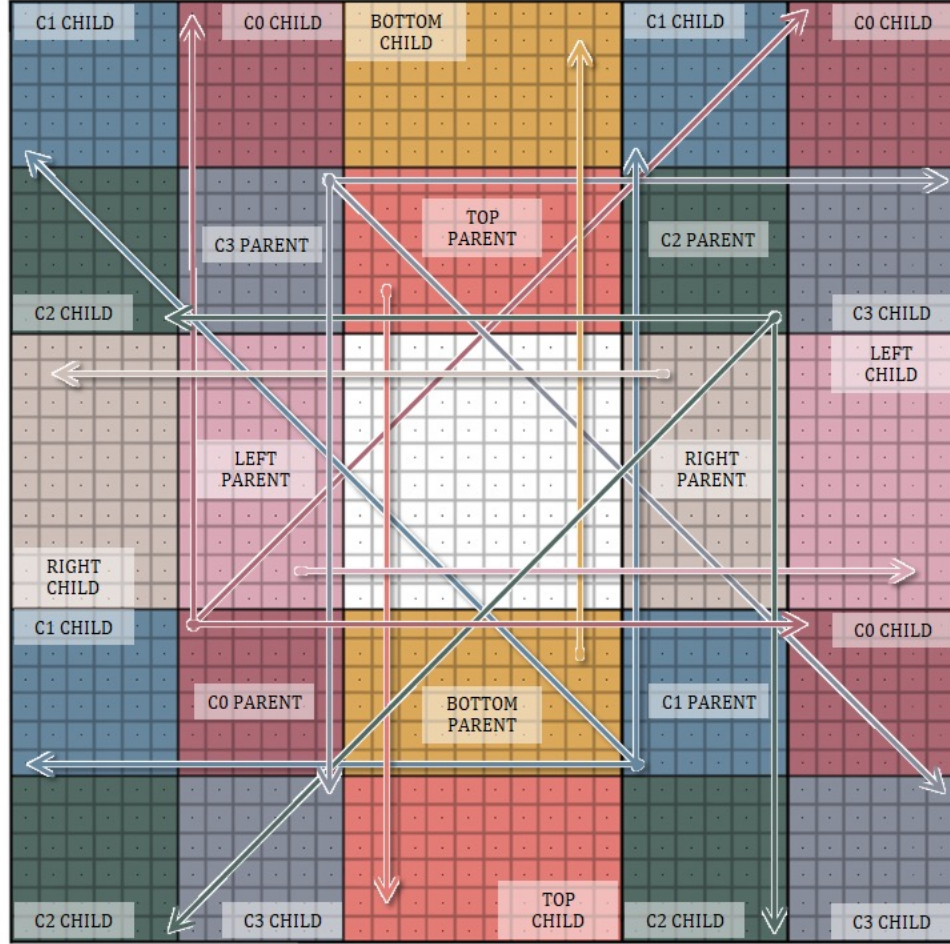
Figure 3.2: Peridynamic particle coupling for periodic boundary condition

Generally, peridynamics requires displacement-based boundary conditions to be applied over a finite volumetric region to ensure accurate enforcement within the domain. As illustrated in Figure 3.2, this fictitious layer is divided into multiple regions, each assigned an identifier (e.g., C0, C1, C2, etc.) and labeled as *CHILD*. Similarly, the exterior volume of the RVE is partitioned into corresponding regions labeled as *PARENT*, each with matching identifiers.

From here we establish the peridynamic volumetric adaptation of the commonly known PBCs, thus ensures compatibility and equilibrium across opposite faces

$$u_i^c = u_i^p + \bar{\varepsilon}_{ij}(x_j^c - x_j^p) \tag{3.23}$$

here, u is displacement, x is reference location. Superscripts c and p refer to *CHILD* and *PARENT* region, respectively. $\bar{\varepsilon}_{ij}$ is the macroscopic strain tensor. Subscripts follow Einstein summation convention with iteration matching the dimension of RVE.

Essentially, for each *CHILD-PARENT* region pair with matching identifier in Figure 3.2, their respective interior particles follow the periodic relation prescribed in (3.23). Note that Figure 3.2 illustrates the displacement coupling relation for a two-dimensional RVE under peridynamic discretisation. However, the formulation in (3.23) is not limited to two dimensions, and the coupling scheme shown in Figure 3.2 can be naturally extended to a three-dimensional cuboidal RVE. A 3D version of the diagram is not shown here, as the 2D illustration is sufficiently clear and the extension to 3D should be intuitive to the reader.

These boundary conditions enforce periodic displacement fluctuations, such that the displacement difference between corresponding points on opposite faces reflects a uniform macroscopic strain. Under these conditions, it can be shown that the volume-averaged microscopic virtual work equals the macroscopic virtual work, thereby satisfying the Hill-Mandel macro-homogeneity condition. A detailed derivation and discussion of the volumetric boundary conditions within peridynamic micromechanics is provided in [4].

3.6 Computation of Effective Properties

To link microscale behaviour with macroscale constitutive response, it is necessary to extract effective material properties from the solution of boundary value problems on the RVE. This section outlines the procedure for computing the effective stiffness tensor from the microscale fields obtained under prescribed boundary conditions. The resulting homogenised properties characterise the overall mechanical behaviour of the heterogeneous material and serve as input for macroscopic continuum models.

3.6.1 Definition of Effective Stiffness Tensor

The homogenised (macroscopic) stiffness tensor is computed from the microscale response of the RVE subjected to a set of controlled boundary conditions. At the microscale, the local constitutive relation is expressed as

$$\sigma_{ij} = C_{ijkl}\,\varepsilon_{kl}, \tag{3.24}$$

where C_{ijkl} denotes the spatially varying stiffness tensor of the heterogeneous microstructure.

The effective (homogenised) stiffness tensor $\overline{C}_{ijkl}$ is defined as the derivative of the macroscopic stress with respect to the macroscopic strain:

$$\overline{C}_{ijkl} := \frac{\partial \overline{\sigma}_{ij}}{\partial \overline{\varepsilon}_{kl}}. \tag{3.25}$$

It governs the macroscopic constitutive behaviour via the relation

$$\overline{\sigma}_{ij} = \overline{C}_{ijkl}\,\overline{\varepsilon}_{kl}, \tag{3.26}$$

where $\overline{\sigma}_{ij}$ and $\overline{\varepsilon}_{kl}$ are the volume-averaged stress and strain fields over the RVE domain.

3.6.2 Boundary Value Problems and Numerical Procedure

The effective stiffness tensor $\overline{C}_{ijkl}$ is computed by solving a sequence of boundary value problems on the RVE, each corresponding to a prescribed macroscopic strain state. For each case, suitable boundary conditions are applied to enforce the imposed macroscopic strain $\overline{\varepsilon}_{kl}$ on the RVE domain. The resulting microscopic displacement fields are obtained by numerically solving the governing equations of motion at the microscale. From the computed displacement field, the microscopic stress tensor σ_{ij} can be subsequently evaluated using peridynamic-based constitutive formulations.

Following the formulation of Lehoucq and Silling [37], the peridynamic Cauchy

stress tensor is defined for an ordinary state-based material as

$$\boldsymbol{\sigma}(\mathbf{x}) = \frac{1}{2} \int_{\mathcal{H}_{\mathbf{x}}} \left[\mathbf{f}(\mathbf{x}, \mathbf{x}') \otimes (\mathbf{x}' - \mathbf{x}) + \mathbf{f}(\mathbf{x}', \mathbf{x}) \otimes (\mathbf{x} - \mathbf{x}') \right] \, dV_{\mathbf{x}'}, \qquad (3.27)$$

where $\mathbf{f}(\mathbf{x}, \mathbf{x}')$ is the pairwise force density and $\mathcal{H}_{\mathbf{x}}$ is the peridynamic horizon. This expression recovers the classical Cauchy stress in the limit of vanishing horizon [37].

In this work, $\boldsymbol{\sigma}(\mathbf{x})$ is evaluated at all material points after the equilibrium solution for each macroscopic strain case. The homogenised RVE stress is then obtained by volume averaging:

$$\overline{\sigma}_{ij} = \frac{1}{|\Omega_{\mathrm{RVE}}|} \int_{\Omega_{\mathrm{RVE}}} \sigma_{ij} \, \mathrm{d}V. \qquad (3.28)$$

This procedure is repeated for six linearly independent macroscopic strain states in three dimensions. The resulting macroscopic stress-strain pairs form a linear system from which the components of $\overline{C}_{ijkl}$ are extracted. For numerical implementation, the fourth-order tensor is represented in Voigt notation.

This workflow follows the established approach for computing peridynamic stress tensors and homogenised moduli, introduces no additional assumptions, and has been validated against analytical and FEM benchmarks for displacement, strain, and stress fields.

Chapter 4

Numerical Implementation

This chapter presents the numerical implementation developed to simulate peridynamic models and compute effective material properties within a multiscale context. The codebase is written in `C++` and leverages a hybrid parallel strategy using OpenMPI [38] and PETSc [39]. The primary goal is to enable scalable and flexible peridynamic simulations for evaluating the homogenised stiffness of representative volume elements (RVEs).

The implementation follows a three-phase simulation workflow:

- **Preprocessing:** Configuration of simulation parameters, domain decomposition, particle discretisation, family (neighbour) search, surface correction, and preprocessing of model-specific constants. Particle distribution across MPI ranks is designed to balance computational load.

- **Solution:** Assembly and solution of the global sparse system of equations, application of boundary conditions (e.g., periodic via Lagrange multipliers), and evaluation of displacement fields using PETSc's iterative solvers.

- **Postprocessing:** Computation of strain and stress fields under the peridynamic formulation, volume-averaging of field quantities, and extraction of effective stiffness using compliance-based methods. The Eigen library is used for local algebraic operations.

```
┌─────────────────────────────┐
│   Software Initialization   │
│    and Problem Definition   │
└─────────────────────────────┘
              │
              ▼
┌─────────────────────────────┐
│   Discretize RVE Geometry   │
│  and Assign Material States │
└─────────────────────────────┘
              │
              ▼
┌─────────────────────────────┐
│  Peridynamic Preprocessing: │
│ Family Search and Surface Correction │
└─────────────────────────────┘
              │
              ▼
┌─────────────────────────────┐
│   Preallocate and Assemble  │
│     Global Sparse System    │
└─────────────────────────────┘
              │
              ▼
┌─────────────────────────────┐
│       Configure Solver      │
│    and Prepare Load Cases   │
└─────────────────────────────┘
              │
              ▼
┌─────────────────────────────┐
│       For i = 1, ..., 6:    │
│ Apply Periodic Boundary Conditions │◄──┐
│       for Strain State i    │   │
└─────────────────────────────┘   │
              │                   │
              ▼                   │
┌─────────────────────────────┐   │
│      Solve Linear System    │   │
│       Using KSPSolve        │───┘
└─────────────────────────────┘
              │
              ▼
┌─────────────────────────────┐
│       Post-processing:      │
│  Peridynamic Stress Evaluation │
│     and Volume Averaging    │
└─────────────────────────────┘
              │
              ▼
┌─────────────────────────────┐
│      Assemble C̄_{ijkl}      │
│      and Export Results     │
└─────────────────────────────┘
```

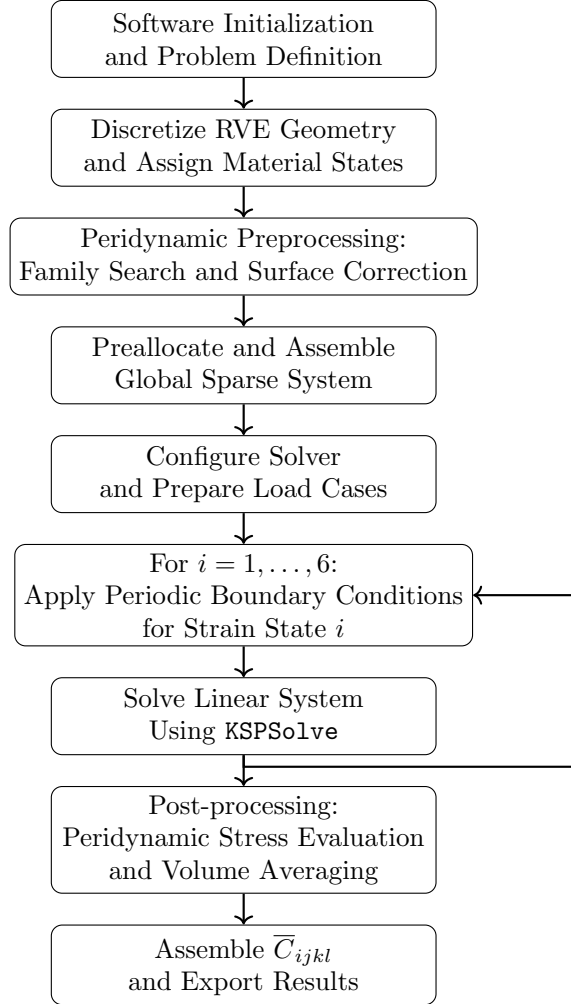Figure 4.1: Computational workflow of the peridynamic homogenisation framework. The procedure includes preprocessing of the RVE, assembly and preallocation of the global sparse system, solver configuration, sequential solution of six periodic boundary value problems corresponding to independent macroscopic strain states, and post-processing to obtain the homogenised stiffness tensor $\overline{C}_{ijkl}$.

The overall numerical workflow of the proposed homogenisation framework is summarised in Figure 4.1. This architecture emphasises modularity and transparency, making it suitable for both high-performance production runs and method development.

4.1 Parallelisation and Domain Decomposition

Parallelism is achieved through MPI-based domain decomposition, where each MPI rank owns a spatial subdomain and manages all particles and interactions contained within it. PETSc provides distributed linear algebra capabilities via the `MATMPIAIJ` and `VECMPI` formats, which automatically handle inter-rank communication during matrix assembly and linear system solutions.

In addition to PETSc-managed objects, the simulation requires the storage and exchange of particle-level data that lie outside PETSc's abstraction (e.g., per-particle states, bond information, or damage-related attributes). These data structures often take the form of nested `C++` containers such as `std::vector<std::vector>`. Because these are not natively supported by PETSc, they are managed using custom MPI communication layers.

During the preprocessing phase, load-balancing information (specifically, the number of particles owned by each rank) is synchronised using `MPI_Allgather`. This information is used to construct a global load-balancing sheet that records the offsets and counts required for later communication routines.

Each MPI rank maintains local copies of the simulation-specific particle data within dedicated communication-handler classes. When synchronisation is required, these classes serialise the nested container structures into one-dimensional buffers based on the previously constructed load-balancing sheet. The flattened buffers are then exchanged across all ranks using `MPI_Allgatherv`, which accommodates variable send and receive sizes. After communication completes, the buffers are deserialised back into their original nested structure to permit local computation.

This communication pattern is applied only to native particle-level data that require nonlocal information stored on remote ranks. Synchronisation is invoked judiciously,

typically deferred until all modifications to the relevant data are completed to minimise communication overhead. A synchronisation step must occur any time a remote fetch is issued on data that may have changed, ensuring consistency across subdomains while avoiding unnecessary MPI traffic.

Because these communication handlers temporarily store complete local copies of intermediate data arrays, memory consumption can become significant. Although this typically has negligible impact on performance for mid-scale problems, explicit memory management becomes essential for large-scale, multi-node simulations. All temporary buffers and auxiliary class data are manually released once they are no longer needed. Without this precaution, the increased memory footprint would hinder the scalability of very large problems. As demonstrated in Section 6.2, the software exhibits near-linear strong scaling for systems exceeding 1.3 million degrees of freedom, a performance trend that would be unattainable without careful control of memory usage.

4.2   Matrix Assembly

The global sparse stiffness matrix is assembled in parallel at the beginning of each loading case. Each MPI rank is responsible for the peridynamic interactions within its local subdomain and performs batched sparse insertions into the distributed PETSc matrix using the `MatSetValues` routine. To ensure scalability, an explicit memory preallocation step is performed before any insertion occurs. A custom preprocessing module estimates the row-wise sparsity pattern by scanning the family structure of each particle and generating approximate nonzero counts for both diagonal and off-diagonal blocks. The estimates intentionally over-predict the number of nonzeros to prevent dynamic memory reallocation during assembly.

Memory preallocation is performed through `MatMPIAIJSetPreallocation`, and the efficiency of the predicted sparsity pattern is monitored using `MatGetInfo`. In practice, the current implementation achieves approximately 110% preallocation relative to the final matrix structure. This slight overestimation ensures that no expensive reallocations occur during matrix insertion, which is essential for high-performance computing environments because insufficient preallocation leads to severe slowdowns and can ren-

der large-scale simulations impractical.

The equations corresponding to real particles inside the RVE are assembled following standard peridynamic discretisation procedures. Because ordinary state-based peridynamics generally leads to a non-symmetric operator, and because periodic boundary conditions are enforced through Lagrange multipliers, no symmetry assumptions are imposed on the global system. Consequently, the stiffness matrix is stored in PETSc's `MATMPIAIJ` format, which supports scalable parallel assembly and efficient SpMV operations on distributed-memory architectures.

Periodic boundary conditions are enforced through Lagrange multiplier constraints, which occupy dedicated rows in the global matrix. These rows correspond to fictitious particles associated with constraint equations, rather than physical particles in the RVE. No physical balance equations are assembled for these fictitious entities; instead, their matrix rows are populated with the coupling terms that enforce displacement periodicity. Only these constraint-related blocks change between loading cases, while the bulk peridynamic contributions remain fixed across simulations. This separation significantly reduces assembly cost during multiscale homogenisation, where multiple boundary value problems with distinct macroscopic strain states must be solved.

4.3   Solver Configuration and Solution Method

The linear systems arising from peridynamic discretisation and periodic constraint enforcement are solved using PETSc's Krylov Subspace Solver (KSP) framework. After the global sparse matrix and right-hand-side vector are assembled, the solution is obtained by issuing a call to the `KSPSolve` routine. This routine performs all iterative updates, residual monitoring, and convergence checks according to the configuration specified by the user.

In this work, the iterative solver is configured to use the `KSPBCGSL` method [40–42], which is a stabilised variant of the BiCG family known for robustness on non-symmetric systems. For all boundary value problems solved during homogenisation, a relative KSP tolerance of 10^{-10} is prescribed. Solver performance was tested across a range of tolerances (from 10^{-8} to 10^{-14}), with negligible differences observed in the results.

The selected value of 10^{-10} represents a reliably achievable and conventionally strict benchmark.

Default PETSc options are used for preconditioning and Krylov subspace parameters. Extensive testing with alternative preconditioners and user-defined configurations did not provide consistent improvements for the types of non-symmetric saddle-point systems encountered in ordinary state-based peridynamics under a homogenisation workflow with periodic boundary conditions. Solver diagnostics, including the number of iterations and the convergence reason, are collected through the `KSPGetIterationNumber` and `KSPGetConvergedReason` routines to provide quantitative verification. No fallback strategies or adaptive restarts were required in production runs.

The solution phase is performed sequentially for the six canonical macroscopic strain states required in first-order computational homogenisation. For each prescribed strain state, periodic boundary conditions are enforced, and the corresponding distributed linear system is assembled across all MPI ranks. A collective call to `KSPSolve` then invokes the parallel iterative solver, which manages all inter-rank communication internally. This process is repeated for all six canonical strains to obtain the complete set of microscopic displacement fluctuation fields needed to compute the homogenised stiffness tensor.

4.4 Summary

The framework outlined in this chapter combines PETSc's high-performance solver backend with custom MPI communication routines to provide a robust and scalable environment for peridynamic simulations. The modular architecture supports a range of model types and boundary conditions and has been validated on HPC platforms with millions of degrees of freedom. This provides a solid foundation for the numerical performance and verification studies presented in subsequent chapters.

Chapter 5

Numerical Results in 2D

5.1 Overview and Setup

This chapter presents a sequence of numerical benchmarks designed to assess the accuracy, robustness, and applicability of the proposed peridynamic homogenisation framework in two dimensions under plane strain assumptions. The study systematically investigates the effective transverse elastic behaviour of heterogeneous composites and the influence of matrix cracking on these properties.

The simulations focus on two main composite types: (i) fibre-reinforced composites with transversely isotropic microstructures, and (ii) pure matrix systems with embedded cracks. All benchmarks assume a plane strain configuration, which is appropriate for modelling thick composite plies and aligns with conventional assumptions in homogenisation theory.

A convergence study based on a boron-aluminium RVE is first conducted to validate the discretisation quality and support the selection of numerical parameters. This is followed by four benchmarks:

- **Benchmark I:** Evaluation of effective transverse elastic properties of a boron-aluminium composite with a fixed fibre volume fraction ($v_f = 0.47$).

- **Benchmark II:** Variation of the effective transverse modulus in glass fibre-epoxy composites across a range of fibre volume fractions, with comparisons to analytical bounds and micromechanical estimates.

- **Benchmark III:** Degradation of effective transverse modulus due to matrix cracking in RVEs, including isolated and coalescing crack configurations.

- **Benchmark IV:** Visualisation of local stress fields in a tiled composite RVE with embedded cracks to examine internal stress distributions and interactions.

All simulations are performed on periodic representative volume elements (RVEs) using the ordinary state-based peridynamic formulation. Bond stiffnesses are assigned using the *Harmonic Mean of Engineering Constants* method (see section 2.5.2) to ensure smooth transitions at material interfaces. A fixed horizon size of $\delta = 3h$ is used, where h is the particle spacing.

The constituent material properties are listed in Table 5.1. Periodic boundary conditions are applied as described in subsection 3.5.3, ensuring compatibility with classical homogenisation theory and enabling the extraction of effective stiffness properties from RVE-level responses.

Table 5.1: Material properties used in 2D benchmarks.

Material	Young's modulus E (GPa)	Poisson's ratio ν
Boron	379.30	0.10
Aluminum	68.30	0.30
Glass	72.40	0.20
Epoxy	2.76	0.35

5.2 Convergence Study

To establish an appropriate discretisation resolution and verify numerical stability, a convergence study is conducted on a single boron-aluminium RVE, consistent with the configuration used in Benchmark I. The goal is to identify the minimum resolution required to obtain accurate effective stiffness results without incurring excessive computational cost.

The RVE geometry used for this study is illustrated in Figure 5.1. It consists of a single cylindrical boron fibre embedded in an aluminium matrix, arranged within a square domain and discretised using a uniform particle grid. Periodic boundary condi-

tions are applied to ensure compatibility with the homogenisation framework described in Chapter 3.


Figure 5.1: Geometry of the 2D boron-aluminium RVE used in the convergence study. The RVE has unit size and contains a single boron fibre embedded in an aluminium matrix, corresponding to a fibre volume fraction of 0.47. Periodic boundary conditions are applied.

Convergence is assessed based on the relative error in three key effective properties: the Young's modulus E_{xx}, the shear modulus G_{xy}, and the Poisson's ratio ν_{xy}. For each property, the relative error is defined as:

$$\text{Relative Error} = \frac{|Q_h - Q_\text{ref}|}{Q_\text{ref}},$$

where Q_h denotes the computed property at grid resolution h, and Q_ref represents the value obtained at the highest resolution used in this study, treated as the reference solution.

Five discretisation levels are considered: 50×50, 100×100, 150×150, 200×200, and 250×250. The results, plotted in Figure 5.2, show a clear monotonic convergence of the computed effective properties with increasing resolution.

To quantify convergence, the relative error of a solution at resolution $N \times N$ is defined with respect to the finest available mesh (250×250). The convergence plot shows that the relative error for all three properties falls below 1% at a resolution of 150×150 and reduces to less than 0.5% at 200×200. More importantly, the rate of decrease in error diminishes significantly at higher resolutions—the improvement from 200×200 to 250×250 is markedly smaller than from 150×150 to 200×200. This indicates the solutions are approaching a mesh-independent asymptotic value, and the 250×250 case can be considered a sufficiently accurate reference for error estimation.

Extrapolating this trend suggests that a hypothetical 270×270 simulation would yield results differing from the 250×250 case by less than 0.1%, which is negligible for the purposes of this study. Therefore, a resolution of 200×200 is selected for all subsequent 2D simulations. This choice provides a practical balance between numerical accuracy and computational efficiency, ensuring results are well within an acceptable error tolerance ($\leq 0.5\%$) relative to the asymptotic converged solution.
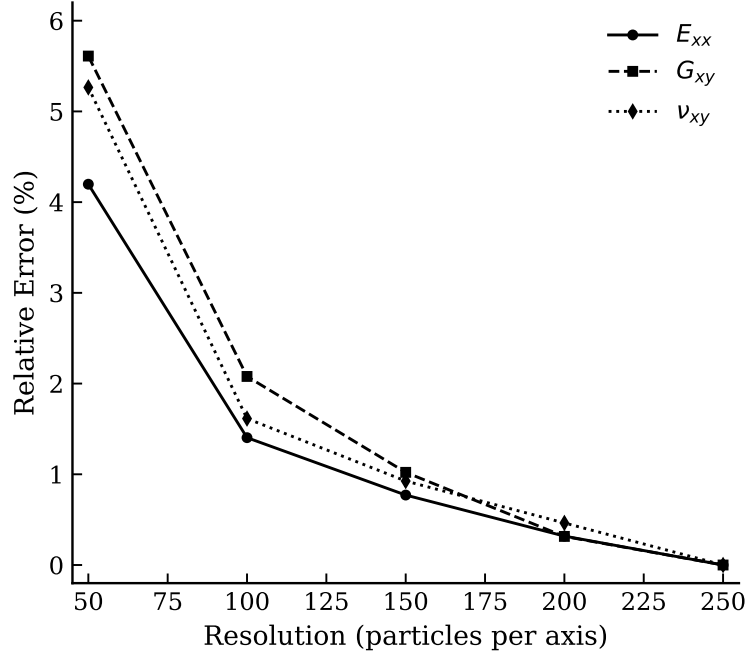

Figure 5.2: Convergence of effective elastic properties in plane strain. Relative errors (%) in E_{xx}, G_{xy}, and ν_{xy} are plotted against increasing RVE resolution.

5.3   Benchmark I: Effective Transverse Elastic Properties of Boron-Aluminium Composite

This benchmark assesses the accuracy of the proposed peridynamic homogenisation framework in predicting the effective in-plane elastic properties of a fibre-reinforced composite under plane strain conditions. The representative volume element (RVE) consists of boron fibres embedded in an aluminium matrix, with a fibre volume fraction of $V_f = 0.47$.

The discretisation used in this benchmark corresponds to the converged resolution identified in section 5.2, and further refinement was found to have negligible influence on the reported results.

Figure 5.3 illustrates the RVE discretisation used in this study. Material points are uniformly distributed, with particles assigned either to the fibre or matrix phase based on their radial position. The particle resolution is 200×200, and periodic boundary conditions are applied following the framework introduced in Chapter 3.


Figure 5.3: Discretisation of the unit-sized composite RVE used in Benchmark I. Each point represents a material point in the peridynamic grid. Colours indicate the material phase: boron fibre (dark grey) and aluminium matrix (light grey).

Table 5.2 presents the computed effective properties from the current peridynamic model, alongside results from established homogenisation methods and available experimental data. The predicted values of E_{xx}, G_{xy}, and ν_{xy} exhibit strong agreement with reference models, generally falling within the range of reported numerical and theoretical estimates.

The slightly elevated stiffness predictions, relative to finite element and asymptotic homogenisation methods, may be partially explained by the plane strain assumption adopted in the present model. In contrast to 3D formulations used in the reference studies, the plane strain constraint suppresses out-of-plane deformation and can result in higher apparent in-plane stiffness. Additional factors, such as differences in how interfacial regions are represented or the use of volumetric periodic boundary conditions may also contribute, though their specific impact is difficult to isolate without further study. Nonetheless, the overall consistency of the results affirms the reliability of the proposed approach in capturing the effective behaviour of fibre-reinforced composites.

Table 5.2: Comparison of effective elastic properties under plane strain (fibre volume fraction = 0.47).

Source	E_{xx} [GPa]	G_{xy} [GPa]	ν_{xy}
Present Model	149.7	46.61	0.282
Madenci et al. [3]	144.4	45.88	0.251
Sun and Vaidya [43]	144.0	45.90	0.290
Yu and Tang [44]	144.1	45.92	0.255
Aboudi et al. [45]	144.0	45.83	0.255
Chamis [46]	156.0	43.6	0.31
Kenaga et al. [47] (Exp.)	140.0	-	-

Notes: Exp. = Experimental.

5.4 Benchmark II: Effective Transverse Elastic Modulus of Glass Fibre-Epoxy Composites Across Varying Fibre Volume Fractions

This benchmark investigates the ability of the proposed peridynamic homogenisation framework to capture the effective transverse stiffness of a unidirectional glass fibre-

epoxy composite across a range of fibre volume fractions. The goal is to assess the consistency of the model when subjected to increasing microstructural anisotropy and phase contrast.

The discretisation used in this benchmark corresponds to the converged resolution identified in section 5.2, and further refinement was found to have negligible influence on the reported results.

The transverse effective modulus E_{22} is computed for fibre volume fractions ranging from 0 to 0.60 in increments of 0.05, using 2D plane strain simulations. The microstructure is modelled as a periodic array of circular glass fibres embedded in an epoxy matrix, with both phases assumed isotropic. For each configuration, the homogenised transverse stiffness is extracted using the peridynamic framework, and the results are compared against several classical micromechanical models.

Figure 5.4 shows the evolution of the effective transverse modulus as a function of fibre volume fraction. The reference models include:

- **Inverse Rule of Mixtures (IROM)**: often associated with the Reuss bound, representing a lower-bound estimate;

- **Halpin-Tsai model**: an empirical semi-theoretical model widely used for fibre-reinforced composites;

- **Modified Mori-Tanaka model**: specifically the transversely isotropic formulation proposed by Tandon and Weng.

Despite being formulated under a plane strain assumption, the peridynamic predictions exhibit strong qualitative agreement with 3D analytical models. The E_{22} values lie slightly above the reference curves, as expected due to the added constraint on out-of-plane deformation. While a direct quantitative comparison is not strictly valid, especially at higher fibre volume fractions, the peridynamic trend closely follows the Halpin-Tsai and Tandon-Weng models, particularly at lower volume fractions.

Across the full range of fibre content, the peridynamic results remain above the IROM (Reuss) lower limit and closely follow established micromechanical models, indicating consistent and physically meaningful behaviour. The close alignment with the

Halpin-Tsai model across most of the domain, and the reasonable proximity to the Tandon-Weng predictions, further support the capability of the proposed homogenisation framework to capture transverse stiffness evolution under simplified 2D assumptions.
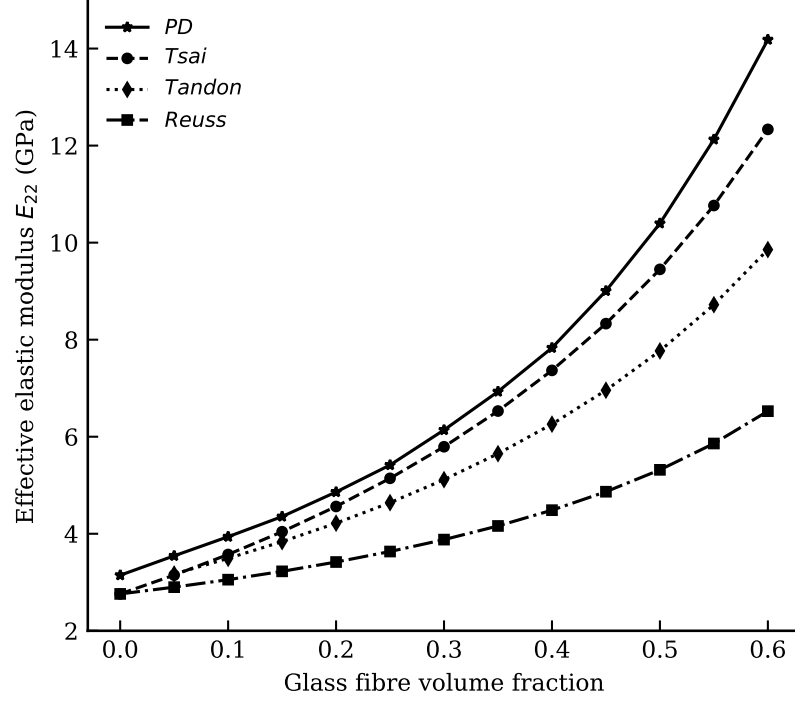

Figure 5.4: Effective transverse Young's modulus E_{22} versus fibre volume fraction V_f. Results from the present peridynamic model (plane strain) are compared with analytical estimates from the inverse rule of mixtures (Reuss [28]), Halpin-Tsai model [36], and modified Mori-Tanaka model (Tandon and Weng [35]).

5.5 Benchmark III: Effective Transverse Modulus Degradation in Cracked Matrix RVEs

This benchmark evaluates the degradation of effective stiffness in cracked matrix materials using the proposed peridynamic homogenisation approach. Two types of representative volume elements (RVEs) are examined, both composed of a homogeneous isotropic matrix material with Young's modulus $E = 2 \times 10^9$ Pa and Poisson's ratio $\nu = 0.3$. The RVEs are unit-sized square domains, subjected to a series of simulations

with increasing crack lengths.

The first RVE contains a single central horizontal crack. The crack length L is swept from 0 to 0.9 in increments of 0.01. The second RVE contains two coalescing horizontal cracks that propagate from their respective origins. For this case, the total crack length is again varied from 0 to 0.9 with a base increment of 0.01, but is locally refined to 0.001 increments in the interval $L \in [0.45, 0.5)$ to accurately resolve the near-singular behaviour as the two cracks approach coalescence.

Two RVE configurations are considered: one containing a single central horizontal crack, and another featuring two coalescing horizontal cracks that propagate from their respective origins. In both cases, the total crack length L is varied from 0 to 0.9. A base increment of 0.01 is used throughout, except for the coalescing crack case, where the increment is refined to 0.001 within the interval $L \in [0.45, 0.5)$ to resolve the steep stiffness transition near crack merging. The upper limit of $L = 0.9$ is intentionally chosen to avoid the singular behaviour that arises when the crack spans the full width of the unit cell ($L = 1$), at which point the RVE becomes discontinuous and the homogenisation problem ill-posed.

The discretisation used in this benchmark corresponds to the converged resolution identified in section 5.2, and further refinement was found to have negligible influence on the reported results.

(a) Single central horizontal crack RVE.

(b) Two coalescing horizontal cracks RVE.

Figure 5.5: RVE configurations used in Benchmark III. Each RVE is of unit size, and the total crack length is L. (a) Single central horizontal crack. (b) Two coalescing horizontal cracks.

Figure 5.5 illustrates the geometry of both RVE configurations, and Figure 5.6 presents the computed transverse stiffness component C_{2222} as a function of crack length. Peridynamic results are compared against analytical solutions (see [48]) drawn from a published model that explicitly accounts for both isolated and coalescing crack configurations. The results show excellent agreement between the peridynamic predictions and the analytical solutions, including the pronounced reduction in stiffness as the cracks grow and merge. In particular, the PD model successfully captures the steep decline in effective stiffness near $L \to 0.5$ in the coalescing case, which is a physically significant feature associated with topological disconnection and loss of load transfer capacity.

To further investigate the mechanical degradation, effective material properties were extracted from the peridynamic simulations and normalised against the undamaged reference case (denoted with an asterisk). Figure 5.7, Figure 5.8, and Figure 5.9 show the evolution of E_{yy}/E_{yy}^*, G_{xy}/G_{xy}^*, and ν_{yx}/ν_{yx}^*, respectively. As expected, all three properties decrease with increasing crack length, reflecting the progressive loss of structural

Figure 5.6: Transverse stiffness component C_{2222} versus crack length L for both single and coalescing crack configurations. Peridynamic results are compared with analytical solutions. The abrupt stiffness reduction near crack coalescence is accurately captured; however, the step-like pattern in the curve arises from the finite discretisation resolution (200×200), rather than a physical discontinuity.

integrity. In contrast, ν_{xy} (calculated as $-S_{12}/S_{11}$) remains nearly constant, while ν_{yx} (as $-S_{21}/S_{22}$) exhibits a notable decline. This anisotropic behaviour arises from the directionality of the crack damage and highlights the sensitivity of the peridynamic method to directional stiffness loss.
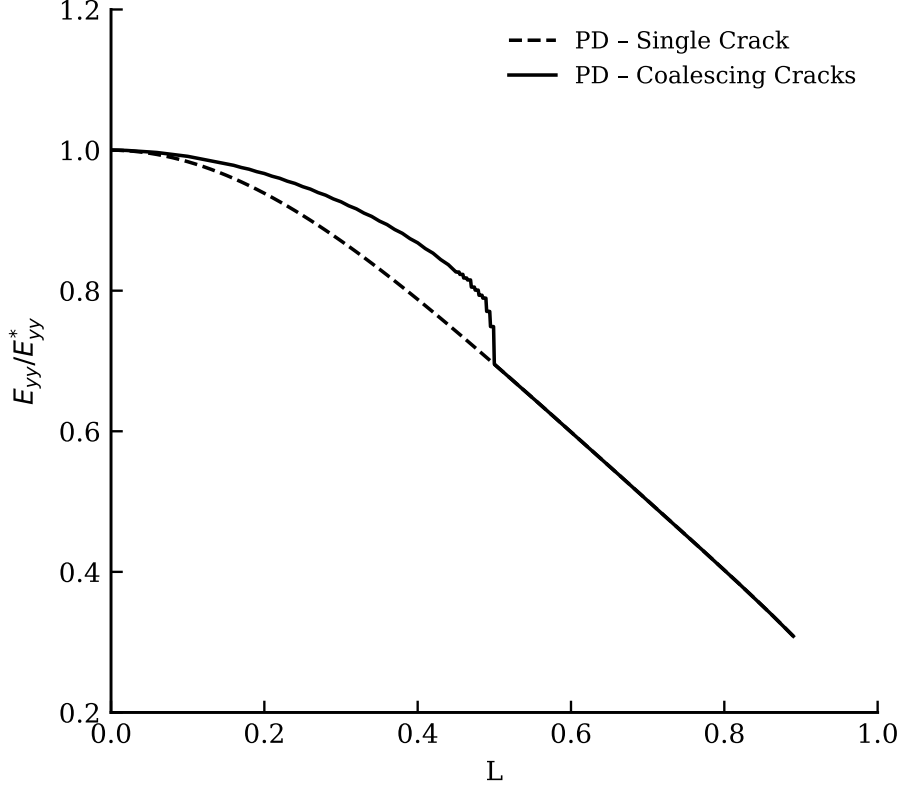
Figure 5.7: Normalised effective Young's modulus E_{yy}/E_{yy}^* as a function of crack length L for single and coalescing crack configurations. Progressive degradation is observed with increasing crack length, consistent with loss of transverse load-carrying capacity.

Because the dominant stiffness degradation occurs in the transverse (y) direction, the comparison against the analytical prediction of C_{2222} provides a quantitative validation of the homogenised response. The remaining effective elastic constants exhibit consistent qualitative trends that reflect the directional influence of matrix cracking. Taken together, these results indicate that the proposed peridynamic homogenisation framework reliably captures both the magnitude and anisotropy of stiffness degradation in cracked media.

Figure 5.8: Normalised effective shear modulus G_{xy}/G_{xy}^* versus crack length L. Both RVE configurations show shear degradation as crack length increases.

Figure 5.9: Normalised Poisson's ratio ν_{yx}/ν_{yx}^* as a function of crack length L. The steady decline in ν_{yx} suggests that vertical transverse deformation becomes less coupled to horizontal loading as the crack length increases, whereas ν_{xy} (not shown) remains relatively unchanged, consistent with the orientation of the cracks.

5.6 Benchmark IV: Stress Field Visualisation in a Cracked Composite RVE (2D)

To examine local stress distributions and interaction phenomena in a complex microstructure, we constructed a 3×3 tiled RVE based on the same boron-aluminium composite used in Benchmark I ($v_f = 0.47$). Each tile has edge length $\ell = \frac{1}{3}$, giving the full RVE a unit size of 1×1. Within each tile, the central fibre remains intact, but two types of microstructural damage are introduced: horizontal fibre-matrix debonding cracks at the interface and vertical matrix cracks embedded symmetrically. All cracks have a uniform length of 0.1ℓ and are modelled as traction-free internal discontinuities without volume, consistent with the methodology outlined in earlier chapters.

The resulting composite RVE structure is shown in Figure 5.10, where both the tiling and symmetric crack layout are apparent. Periodic boundary conditions are imposed on all sides, and standard strain-controlled homogenisation procedures are applied, following the same workflow used in the prior benchmarks.

Figure 5.10: RVE configuration used in Benchmark IV. The unit-sized composite RVE is formed by tiling the boron-aluminium microstructure used in Benchmark I into a 3×3 grid. Debonding cracks are placed at fibre interfaces, and matrix cracks are embedded vertically. All cracks are of length 0.1ℓ. Periodic boundary conditions are applied.

For each of the three macroscopic strain loading modes (ε_{xx}, ε_{yy}, and ε_{xy}), the corresponding dominant stress component (σ_{xx}, σ_{yy}, and σ_{xy}, respectively) is extracted and visualised across the RVE. Additionally, the von Mises stress field is computed and plotted for each loading case to capture combined stress concentrations. In total, six field plots are presented, grouped by loading condition.

Stress Field: σ_{xx} under Macroscopic Strain ε_{xx}

Stress Field: σ_{vm} under Macroscopic Strain ε_{xx}

(a) σ_{xx}

(b) σ_{vm}

Figure 5.11: Stress fields under macroscopic strain ε_{xx}. (a) Stress component σ_{xx}. (b) von Mises stress σ_{vm}.

Stress Field: σ_{yy} under Macroscopic Strain ε_{yy}

Stress Field: σ_{vm} under Macroscopic Strain ε_{yy}

(a) σ_{yy}

(b) σ_{vm}

Figure 5.12: Stress fields under macroscopic strain ε_{yy}. (a) Stress component σ_{yy}. (b) von Mises stress σ_{vm}.
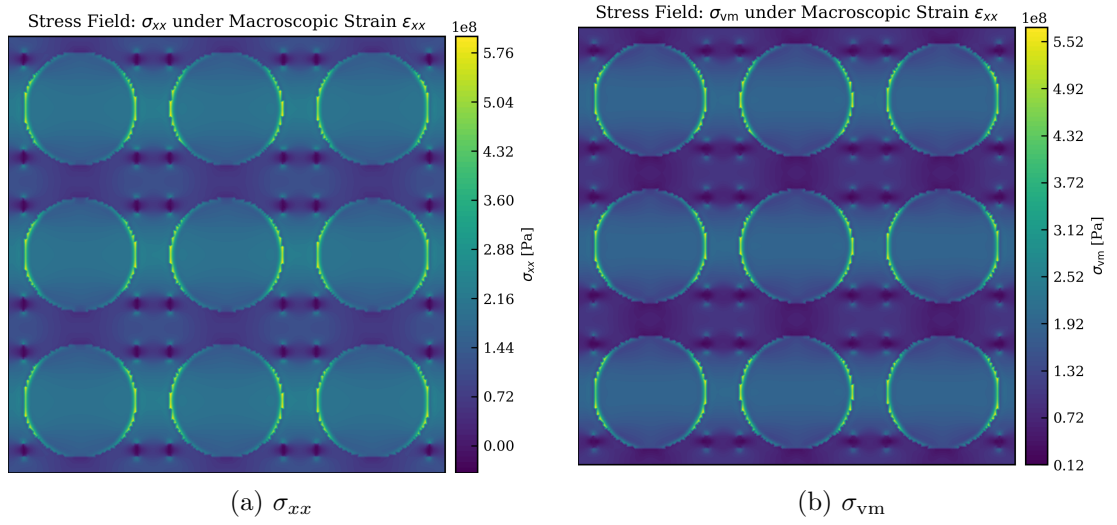
(a) σ_{xy} (b) σ_{vm}

Figure 5.13: Stress fields under macroscopic shear strain ε_{xy}. (a) Stress component σ_{xy}. (b) von Mises stress σ_{vm}.
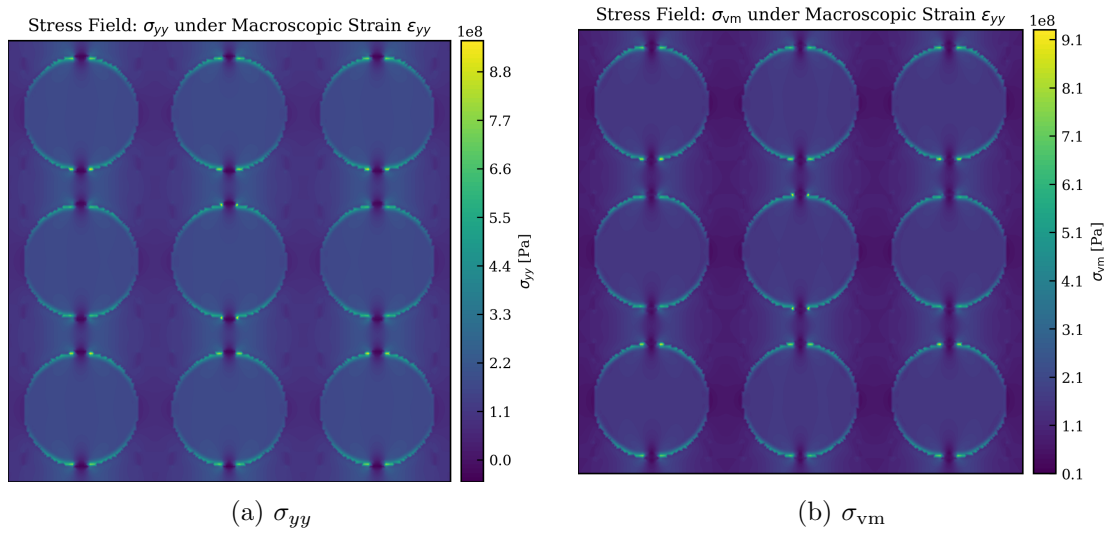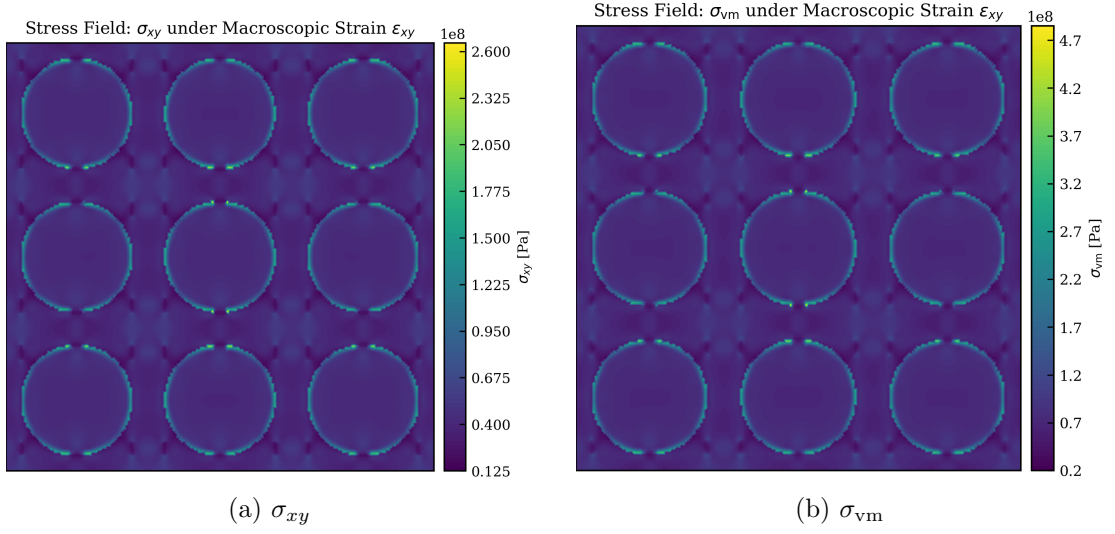
As expected, the stress fields display strong periodicity consistent with the imposed boundary conditions. Stress concentrations are observed near fibre-matrix interfaces, particularly at debonding regions, and at crack tips where local gradients are amplified. Regions where cracks are in close proximity exhibit clear interaction effects, often resulting in intensified local stress bands. The von Mises fields further highlight these high-stress zones, revealing how the interplay between microstructural heterogeneity and damage governs local stress amplification.

All plots are presented in physical units (Pa), without normalisation, to preserve the true scale of stress variation. While no direct analytical comparison is provided for this benchmark, the observed field patterns and concentration zones are qualitatively consistent with results reported in the literature, such as in [49, 50]. In particular, similar stress localisation and crack interaction behaviours have been noted in prior micromechanical studies, lending confidence to the fidelity of the present approach, even though the material systems and geometric configurations differ.

Chapter Summary

This chapter presented a series of two-dimensional numerical benchmarks designed to assess the predictive capabilities of the proposed peridynamic homogenisation framework. Through comparisons with analytical solutions, classical micromechanics models, and detailed field visualisations, the method demonstrated strong consistency and physical fidelity across a range of composite configurations and damage scenarios. In particular, the effective stiffness predictions aligned well with theoretical models at various fibre volume fractions, and the degradation behaviour in cracked matrices was accurately captured. The qualitative stress field plots further revealed the localised effects of damage and interface discontinuities.

These 2D results form the basis for the extended 3D analyses presented in the next chapter, where the framework is applied to three-dimensional RVEs to evaluate its scalability and robustness in capturing complex material behaviour.

Chapter 6

Numerical Results in 3D

6.1 Overview and Setup

This chapter extends the computational homogenisation framework developed in chapter 5 to three-dimensional composite materials. Four benchmarks are presented, progressively addressing stiffness prediction in fibre composites, particulate composites, crack-induced property degradation, and local stress field visualisation in damaged microstructures. The purpose of these benchmarks is to demonstrate the method's applicability in predicting effective anisotropic properties, modelling damage-driven stiffness loss, and resolving local stress interactions within fully three-dimensional representative volume elements (RVEs).

The material systems analysed include both boron-aluminium and glass-epoxy composites, representing typical high-stiffness and particulate-reinforced materials, respectively. Their linear elastic properties are summarised in Table 6.1.

Table 6.1: Material properties used in 3D benchmarks.

Material	Young's Modulus (GPa)	Poisson's Ratio
Boron	379.30	0.10
Aluminium	68.30	0.30
Glass	72.40	0.20
Epoxy	2.76	0.35

All simulations employ uniform grid discretisation with the peridynamic horizon set

at three times the particle spacing. Periodic boundary conditions are enforced on all RVE boundaries to ensure compatibility with homogenisation theory. For modelling material interfaces, harmonic mean sampling using ten quadrature points along each bond is applied consistently across the entire domain. Bonds intersected by cracks are deactivated based on ray-tracing using a modified Möller-Trumbore algorithm with edge-contact handling.

Numerical solutions are obtained using the PETSc library, specifically the KSP-BCGSL iterative solver, with a convergence tolerance of 10^{-10}. All benchmarks are executed on a single high-performance computing node utilising 40 MPI ranks, with each rank allocated 3.5 GB of RAM. Benchmark-specific discretisation resolutions and runtimes are reported individually within the corresponding sections.

The subsequent sections detail the setup, results, and analysis for each benchmark.

6.2   Convergence Study

This section investigates the numerical convergence behaviour of the proposed peridynamic homogenisation framework applied to a three-dimensional boron-aluminium composite RVE with a fibre volume fraction of $v_f = 0.47$ (see Table 6.1 for material parameters). As illustrated in Figure 6.1, the geometry consists of a single cylindrical boron fibre embedded in an aluminium matrix, oriented along the x-axis (corresponding to direction 1 in tensor notation).

Figure 6.1: Geometry of the 3D boron-aluminium RVE used in the convergence study. A single cylindrical boron fibre is embedded in an aluminium matrix and aligned along the x-axis (direction 1).

Eight discretisation levels were tested, corresponding to nominal resolutions of [20, 30, 40, 50, 60, 70, 80, 90] particles per spatial dimension. To enforce periodic boundary conditions, a padding layer is added at each boundary. The thickness of this padding is set to three particles, corresponding exactly to the horizon size of $\delta = 3\Delta x$, ensuring that all interactions within the horizon of points near the domain boundary are fully accounted for across the periodic interface. This results in actual particle counts of [26, 36, 46, 56, 66, 76, 86, 96] along each axis.

To evaluate sensitivity to spatial resolution, the relative error of each homogenised stiffness component is computed with respect to the highest-resolution result (90^3 particles), treated as the reference solution:

$$\text{Relative Error} = \frac{\left| C_{ijkl}^h - C_{ijkl}^{\text{ref}} \right|}{C_{ijkl}^{\text{ref}}},$$

where C_{ijkl}^h is the computed stiffness at resolution level h, and C_{ijkl}^{ref} denotes the corresponding value at the finest resolution.

The results are grouped into three categories:

- Axial moduli: $C_{1111}, C_{2222} = C_{3333}$

- Shear moduli: $C_{2323}, C_{1212} = C_{1313}$

- Coupling moduli: $C_{1122} = C_{1133}, C_{2211} = C_{3311}, C_{2233} = C_{3322}$

Figure 6.2: Convergence of axial stiffness components. Relative error (%) in C_{1111}, C_{2222}, and C_{3333} as a function of RVE resolution.

As shown in Figure 6.2, the tensor component C_{1111} becomes nearly insensitive to resolution beyond 50^3, with relative errors remaining below 0.1%. The components C_{2222} and C_{3333} are identical due to the transverse isotropy of the RVE geometry. The anomalous dip at the 40^3 resolution was investigated for robustness. Multiple simulations with different solver configurations confirmed the result is repeatable and not a numerical artefact. This localized deviation is therefore attributed to a resolution-specific discretisation effect. At this specific particle density, the geometric representation of the fibre-matrix interface may lead to a slight, non-systematic shift in the apparent volume fraction or local connectivity, causing a small deviation from the overall trend.
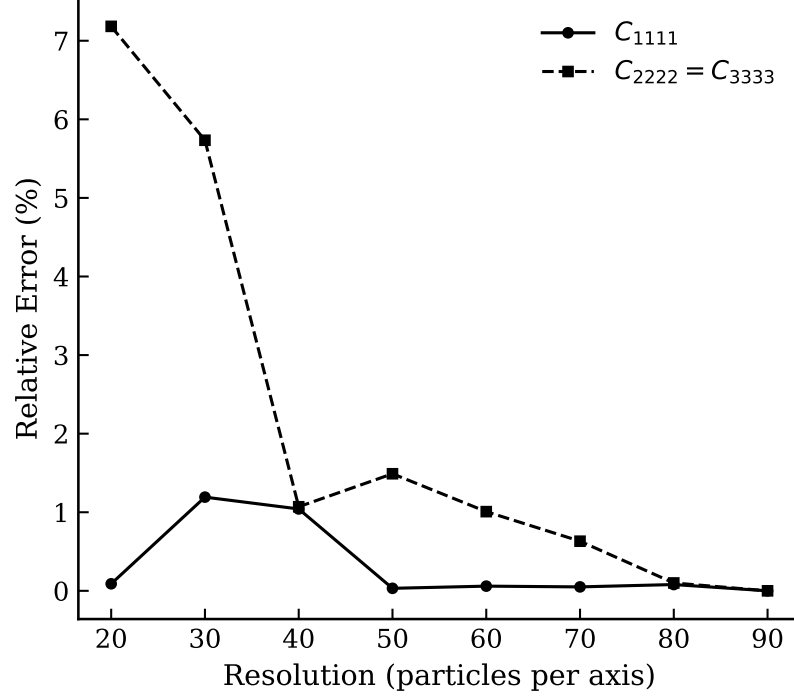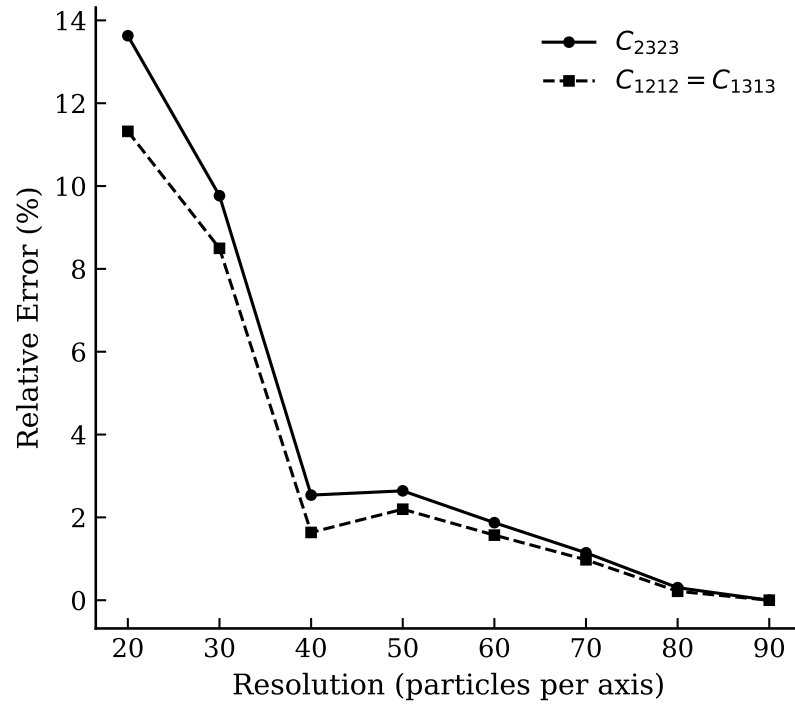
Figure 6.3: Convergence of shear stiffness components. Relative error (%) in C_{1212}, C_{1313}, and C_{2323} as a function of RVE resolution.

As shown in Figure 6.3, the shear stiffness components converge smoothly with increasing resolution. The values of C_{1212} and C_{1313} fall below 1% relative error after the 70^3 case, reflecting rapid convergence. These two components are identical due to the transverse isotropy of the microstructure. In contrast, C_{2323} exhibits slower convergence, likely due to its greater sensitivity to the discretised fibre-matrix interface in the 2-3 plane.

discretisation

Figure 6.4: Convergence of coupling stiffness components. Relative error (%) for C_{1122}, C_{1133}, C_{2233}, C_{3322}, C_{2211}, and C_{3311} plotted against increasing RVE resolution.

In Figure 6.4, the coupling components C_{1122} and C_{1133} exhibit minimal sensitivity to resolution, with relative errors below 0.3% even at coarse discretisations. Cross-directional terms such as C_{2233} and C_{2211} show a more gradual convergence trend, but their errors fall below 1% at resolutions of 70^3 and higher.

Overall, the convergence trends demonstrate that increasing the discretisation resolution improves the accuracy and stability of the computed effective stiffness components. Most tensor entries achieve sub-1% relative error at resolutions above 70^3, with many converging significantly earlier. The anomaly observed at the 40^3 level seen across several components is likely due to geometric representation artefacts, such as abrupt shifts in realised fibre volume fraction resulting from coarse discretisation.

A resolution of 90^3 (corresponding to $96 \times 96 \times 96$ particles and approximately 2.65 million degrees of freedom) is adopted for all subsequent three-dimensional benchmarks, unless otherwise specified. This choice offers high accuracy while remaining within the memory limits of a single-node, 40-core HPC configuration. Higher resolutions, such as

83

100^3, exceed available memory and require multi-node or high-memory setups. Thus, 90^3 represents the optimal balance between fidelity and computational feasibility for the present study.

Figure 6.5: Wall-clock runtime versus total number of degrees of freedom for homogenisation of boron-aluminium composite RVE. The non-uniform spacing of DOF values reflects the cubic growth of the discretisation. The results demonstrate the scalability of the MPI-based peridynamic homogenisation framework with increasing problem size.

In Figure 6.5, we report the wall-clock runtime for the five RVE discretisations executed on a uniform 40-rank configuration. For smaller problem sizes (below approximately 1 million DOF), the runtime departs from linearity because initialisation procedures, family-list construction, and communication overheads are comparable to the computational workload. Once the problem size exceeds about 1.3 million DOF, the runtime demonstrates near-linear growth, indicating that the implementation has reached its asymptotic scaling regime.

6.3 Benchmark V: Effective Stiffness of Boron-Aluminium Composite

This benchmark evaluates the accuracy of the proposed peridynamic homogenisation method by computing the effective stiffness of a three-dimensional boron-aluminium composite. The RVE considered here is identical to that used in the preceding convergence study, with a fibre volume fraction of $v_f = 0.47$. The geometry consists of a single cylindrical boron fibre embedded in an aluminium matrix and aligned along direction 1 (x-axis). The homogenisation is performed using the finest tested resolution of 90^3 interior particles, which, accounting for periodic boundary padding, results in a 96^3 grid and approximately 2.65 million degrees of freedom.

Figure 6.6: Discretisation of the boron-aluminium RVE used in Benchmark V. The fibre is aligned along direction 1 (x-axis), with periodic boundary conditions applied in all directions(not shown).

The homogenised effective stiffness tensor $\boldsymbol{C}^*$ and its inverse, the compliance tensor $\boldsymbol{S}^*$, are obtained by applying macroscopic strain-controlled periodic boundary conditions, as detailed in earlier methodology chapters. The full tensors are presented below.

Chapter 6.   Numerical Results in 3D

All values are in physical units (Pa and Pa^{-1}). Small off-diagonal terms are retained as-is to reflect numerical fidelity and floating-point limits. These terms do not exceed expected tolerances, and the overall symmetry and transverse isotropy of the RVE are well captured.

$$C^* = \begin{bmatrix} 2.31128\text{E}+11 & 4.03574\text{E}+10 & 4.03574\text{E}+10 & 0.0463754 & 0.0564391 & 0.185924 \\ 3.68198\text{E}+10 & 1.64514\text{E}+11 & 4.41003\text{E}+10 & -0.0230047 & 0.265067 & 10.6513 \\ 3.68198\text{E}+10 & 4.41003\text{E}+10 & 1.64514\text{E}+11 & -0.0234455 & 0.08103 & 0.164594 \\ -4.20361 & 14.2165 & -9.69425 & 4.75119\text{E}+10 & -0.405878 & -5.50247 \\ -34.9252 & -0.113261 & 0.430695 & 1.0164\text{E}-05 & 5.62132\text{E}+10 & 0.388198 \\ 12.1984 & 0.171725 & 1.56149 & -5.93308\text{E}-05 & -0.389043 & 5.62132\text{E}+10 \end{bmatrix} \text{Pa}$$

$$S^* = \begin{bmatrix} 4.611\text{E}-12 & -8.920\text{E}-13 & -8.920\text{E}-13 & -5.373\text{E}-24 & 8.625\text{E}-25 & 1.564\text{E}-22 \\ -8.138\text{E}-13 & 6.707\text{E}-12 & -1.598\text{E}-12 & 3.253\text{E}-24 & -2.850\text{E}-23 & -1.263\text{E}-21 \\ -8.138\text{E}-13 & -1.598\text{E}-12 & 6.707\text{E}-12 & 3.330\text{E}-24 & -1.314\text{E}-24 & 2.859\text{E}-22 \\ 4.854\text{E}-22 & -2.412\text{E}-21 & 1.768\text{E}-21 & 2.105\text{E}-11 & 1.520\text{E}-22 & 2.060\text{E}-21 \\ 2.869\text{E}-21 & -5.284\text{E}-22 & -6.088\text{E}-22 & -3.806\text{E}-27 & 1.779\text{E}-11 & -1.229\text{E}-22 \\ -9.755\text{E}-22 & 2.175\text{E}-22 & 1.215\text{E}-23 & 2.222\text{E}-26 & 1.231\text{E}-22 & 1.779\text{E}-11 \end{bmatrix} \text{Pa}^{-1}$$

The key effective moduli extracted from the homogenised stiffness tensor are summarised and compared against various micromechanical models and experimental data in Table 6.2. The proposed method achieves excellent agreement with both finite element and peridynamic unit cell approaches, with deviations generally under 3%. Slight differences in Poisson's ratios may arise due to interface treatment strategies, such as the weighted average approach described in section 2.5.2.

Table 6.2: Comparison of effective elastic properties (fibre volume fraction = 0.47). All stiffness and shear modulus values are in [GPa].

Source	E_1	E_2	G_{12}	G_{23}	ν_{12}	ν_{23}
Present Model	216.88	149.11	56.21	47.51	0.193	0.238
Madenci et al. [3]	215.05	144.4	54.3	45.8	0.195	0.251
Sun and Vaidya [43]	215	144	57.2	45.9	0.19	0.29
Yu and Tang [44]	215.3	144.1	54.39	45.92	0.195	0.255
Aboudi et al. [45]	215.4	144	54.34	45.83	0.195	0.255
Chamis [46]	214.0	156.0	62.6	43.6	0.2	0.31
Kenaga et al. [47] (Exp.)	216.0	140	52.0	-	0.29	-

Notes: Exp. = Experimental.

These results confirm the ability of the proposed framework to deliver high-fidelity predictions for anisotropic composites in three dimensions. The computed stiffness matrix captures both directional moduli and coupling terms in accordance with theoretical expectations, with physically reasonable deviations from other models. These discrepancies are most likely attributable to differences in interface modelling strategies. Overall, the model's agreement with a wide range of literature benchmarks builds confidence in its applicability to more complex field analyses and damage scenarios in later sections.

The full-scale homogenisation was performed on a single high-performance computing node with 40 MPI ranks, each allocated a memory limit of 3.5 GB per core. The total runtime for the 90^3 case was 7 hours, 33 minutes, and 37 seconds. The primary performance bottleneck arises from the interface treatment strategy (*Harmonic Mean of Engineering Constants*) described in section 2.5.2, where each bond is sampled at 10 equidistant points along its length to account for spatial variation in material properties. Bond stiffness is then computed based on peridynamic parameters assembled from the harmonically averaged engineering constants. At present, this sampling is applied uniformly to all bonds, regardless of whether they cross a material interface, which introduces significant computational overhead. While this approach ensures consistency and simplifies implementation, it is recognised as an inefficiency that could be addressed through selective sampling in future optimisation. Nevertheless, the high-

resolution case remains solvable on a single node, demonstrating the robustness of the present formulation under demanding conditions.

6.4 Benchmark VI: Effective Properties of Glass-Epoxy Composites Across Varying Fibre Volume Fractions

This benchmark investigates the predictive capability of the proposed peridynamic homogenisation framework in capturing the effective elastic behaviour of glass-epoxy composites across a range of fibre volume fractions. A simplified representative volume element (RVE) is constructed by embedding a single spherical glass inclusion inside an epoxy matrix, with the inclusion volume fraction v_f varied from 0.0 to 0.5 in increments of 0.1. All RVEs are discretised using the same high-fidelity peridynamic resolution of 90^3 grid points, corresponding to 96^3 material points after periodic boundary padding. The peridynamic horizon is set to three times the particle spacing, and periodic boundary conditions are applied on all faces.

A sample discretisation for the $v_f = 0.3$ case is shown in Figure 6.7, where glass and epoxy particles are colour-coded to illustrate phase distribution. The spherical inclusion is centrally embedded, and the entire structure is constructed to maintain perfect periodicity.

Figure 6.7: Discretisation of the composite RVE used in Benchmark VI, with fibre volume fraction $v_f = 0.3$. A single spherical glass inclusion (dark red) is embedded in an epoxy matrix (silver grey). Periodic boundary conditions are applied on all faces (not shown).

The effective elastic constants are extracted via the standard homogenisation procedure outlined in Chapter 3, with the same numerical settings used throughout: solver tolerance 10^{-10}, high-performance computing (HPC) node with 40 MPI ranks, and memory allocation of 3.5 GB per rank. Each volume fraction case requires approximately 7.5 hours to complete. As discussed in section 6.2, the resolution 90^3 has been validated for similar material systems and is used consistently here to ensure comparability.

Analytical micromechanics models typically express effective composite behaviour in terms of bulk modulus K and shear modulus G, rather than Young's modulus and Poisson's ratio. Accordingly, the present benchmark focuses on these two quantities to enable direct and meaningful comparison across different theoretical estimates. Analytical bounds and classical micromechanics estimates are included for reference: Voigt and Reuss bounds (upper and lower bounds), Hashin-Shtrikman (HS$^+$/HS$^-$) bounds, and Mori-Tanaka (MT) theory.

Figure 6.8: Effective bulk modulus K of glass-epoxy composite versus fibre volume fraction v_f. Comparison is made against Voigt/Reuss bounds, Hashin-Shtrikman bounds, and Mori-Tanaka theory. The PD model prediction tracks HS$^-$ and MT closely.

Figure 6.9: Effective shear modulus G of glass-epoxy composite versus fibre volume fraction v_f. Similar to the bulk modulus, the PD model aligns well with the Hashin-Shtrikman lower bound and Mori-Tanaka estimate.

As shown in Figures 6.8 and 6.9, the peridynamic model yields effective bulk and shear moduli that remain fully bounded by classical estimates across all volume fractions. The Voigt-Reuss envelope provides the widest bounds, with $HS^{\pm}$ offering tighter constraints. In this case, the Mori-Tanaka predictions are observed to coincide with the Hashin-Shtrikman lower bounds across all volume fractions, consistent with expectations for isotropic composites containing well-dispersed spherical inclusions.

The PD results closely follow the HS^- and MT predictions, particularly for $v_f \leq$ 0.35, indicating that the proposed framework is well suited for modelling moderately reinforced composites. Slight upward deviation from the HS^- bound is observed at higher fibre content, which remains physically plausible and may be attributed to local interface averaging effects introduced by the weighted bond sampling strategy.

Overall, the benchmark demonstrates that the peridynamic homogenisation method not only preserves theoretical bounding behaviour, but also delivers quantitatively con-

sistent estimates in comparison to widely accepted analytical models. These results reinforce the method's applicability in isotropic-inclusion systems and motivate further exploration into damage-driven property degradation in the next benchmark.

6.5 Benchmark VII: Effective Property Degradation in Cracked Matrix RVEs

This benchmark examines the degradation of effective stiffness in a cracked matrix RVE, extending the previous two-dimensional analysis into a fully three-dimensional setting. The RVE consists of a pure matrix material with no fibre inclusions. The matrix is modelled as a solid plate with explicit thickness, with Young's modulus $E = 2\,\text{GPa}$ and Poisson's ratio $\nu = 0.3$.

Two crack configurations are considered: a single central crack, and a pair of coalescing horizontal cracks. As illustrated in Figure 6.10, both crack systems are oriented in the xz-plane and span the full thickness of the RVE. Cracks are explicitly represented as zero-thickness surfaces using triangle mesh geometry, allowing precise resolution of surface intersections in the peridynamic bond-breaking procedure.

(a) Single horizontal central crack.

(b) Two coalescing horizontal cracks.

Figure 6.10: Planar x-y view of the cracked RVE used in Benchmark VII. Each RVE is of size $1 \times 1 \times 0.08$, and the total crack length is L. Periodic boundary conditions applied. (a) Single horizontal crack; (b) Two coalescing horizontal cracks.

The RVE is discretised using a uniform grid of $100 \times 100 \times 8$ particles, consistent with prior peridynamic homogenisation studies [3]. Accounting for periodic boundary condition padding, the actual discretisation employed is $106 \times 106 \times 14$. The peridynamic horizon is maintained at three times the particle spacing. Cracks are incorporated using the modified Möller-Trumbore ray-tracing algorithm with edge-contact handling, as described in earlier methodology sections. This approach ensures accurate identification and deactivation of bonds intersected by the crack surfaces.

Simulations are conducted using a Krylov Subspace solver with a tolerance of 10^{-10}, employing 40 MPI ranks distributed across a single high-performance computing node. The total crack length L is varied from 0 to 0.95 in increments of 0.01. For the two-crack coalescence scenario, an additional refinement is applied over $L \in [0.45, 0.5)$ with smaller increments of 0.001 to resolve the rapid stiffness degradation during crack merging. As $L \to 1$, the RVE approaches geometric separation into two disconnected blocks, leading to a rapid decline in effective stiffness. To avoid numerical singularities arising from complete severance at $L = 1$, the upper limit of L is capped at 0.95.

The evolution of the effective stiffness component C_{2222} as a function of total crack length L is presented in Figure 6.11, covering both the single-crack and coalescing-crack configurations. Analytical predictions from the asymptotic solutions of Markenscoff [48] are included for comparison.

Figure 6.11: Effective stiffness component C_{2222} versus total crack length L for both single-crack and coalescing-crack configurations. Peridynamic predictions are compared against analytical solutions [48]. The sharp stiffness reduction during crack coalescence is accurately captured. The step-like pattern in the curves reflects discretisation-induced artefacts from the finite resolution ($100 \times 100 \times 8$), rather than physical discontinuities.

Across both configurations, the peridynamic model successfully reproduces the expected stiffness degradation trends, showing good agreement with analytical predictions. The progressive reduction of C_{2222} is captured accurately, including the abrupt stiffness drop associated with crack coalescence. Minor step-like variations observed in the numerical results are attributed to finite discretisation resolution rather than

physical effects.

Overall, these results confirm the capability of the proposed framework to predict effective property degradation in cracked matrix composites under fully three-dimensional conditions. The method accurately captures both crack propagation and interaction effects, demonstrating robustness for modelling fracture-driven stiffness degradation.

6.6 Benchmark VIII: Stress Field Visualisation in a Cracked Composite RVE (3D)

To examine local stress distributions and crack interaction phenomena in a three-dimensional microstructure, a 3×3 tiled RVE was constructed based on the boron-aluminium composite system used in Benchmark V ($v_f = 0.47$), forming a plate-like RVE of dimensions $0.05 \times 1 \times 1$. Fibres and cracks are extruded along the thickness direction (x). Within each tile, the central boron fibre remains intact, while two types of microstructural damage are introduced: horizontal fibre-matrix debonding cracks located at the fibre interfaces, and vertical matrix microcracks positioned symmetrically between fibres. All cracks have a uniform length of 0.1ℓ and are modelled as internal zero-thickness surfaces, consistent with the methodology outlined in earlier chapters.

A planar view of the RVE configuration in the y-z plane is shown in Figure 6.12, illustrating the tiled structure and crack placement. Each tile has an edge length of $\ell = \frac{1}{3}$. Periodic boundary conditions are applied to all faces, and strain-controlled macroscopic loading is imposed following the homogenisation procedure described previously.

Figure 6.12: Planar y-z view of the composite RVE used in Benchmark VIII. The RVE has size $0.05 \times 1 \times 1$. The fibres and cracks extend fully through the thickness direction (x). Fibre-matrix debonding cracks and matrix microcracks are included as internal zero-thickness surfaces. Periodic boundary conditions applied.

The RVE is discretised using a uniform grid of $10 \times 200 \times 200$ particles along the x, y, and z axes, respectively. Accounting for periodic boundary condition padding, the actual discretisation employed is $16 \times 206 \times 206$. The peridynamic horizon is maintained at three times the particle spacing. Cracks are represented by triangle mesh surfaces and used to define bond-breaking conditions via the modified Möller-Trumbore ray-tracing algorithm. Bond properties across material interfaces are evaluated using harmonic mean sampling with 10 quadrature points along each bond, consistent with prior benchmarks.

Simulations are performed on a single high-performance computing node with 40 MPI ranks, each allocated 3.5 GB of RAM. The total runtime for this benchmark was 7 hours, 15 minutes, and 30 seconds.

Figure 6.13: Discretisation of the composite RVE used in Benchmark VIII. The RVE is discretised using a uniform grid of $10 \times 200 \times 200$. Fibre particles are shown in dark red and the matrix particles in grey. Cracks are not visualised, as their zero-thickness mesh representation does not influence particle positions.

To investigate stress transfer and damage interaction mechanisms, six independent macroscopic strain cases are applied: ε_{xx}, ε_{yy}, ε_{zz}, γ_{yz}, γ_{zx}, and γ_{xy}. For each loading mode, the corresponding dominant stress component (σ_{xx}, σ_{yy}, σ_{zz}, σ_{yz}, σ_{zx}, and σ_{xy}, respectively) and the von Mises stress field are extracted. All stress fields are evaluated on the mid-plane slice in the y-z plane, corresponding to the mid-thickness in the x-direction. Field visualisation is performed using triangular contour interpolation with 100 levels.

Figures 6.14 through 6.19 present the stress fields for all six loading cases.

(a) σ_{xx} (b) von Mises

Figure 6.14: Stress fields at the mid-plane of the RVE under macroscopic strain ε_{xx}. (a) Stress component σ_{xx}. (b) von Mises stress σ_{vm}.

(a) σ_{yy} (b) von Mises

Figure 6.15: Stress fields at the mid-plane of the RVE under macroscopic strain ε_{yy}. (a) Stress component σ_{yy}. (b) von Mises stress σ_{vm}.

(a) σ_{zz}
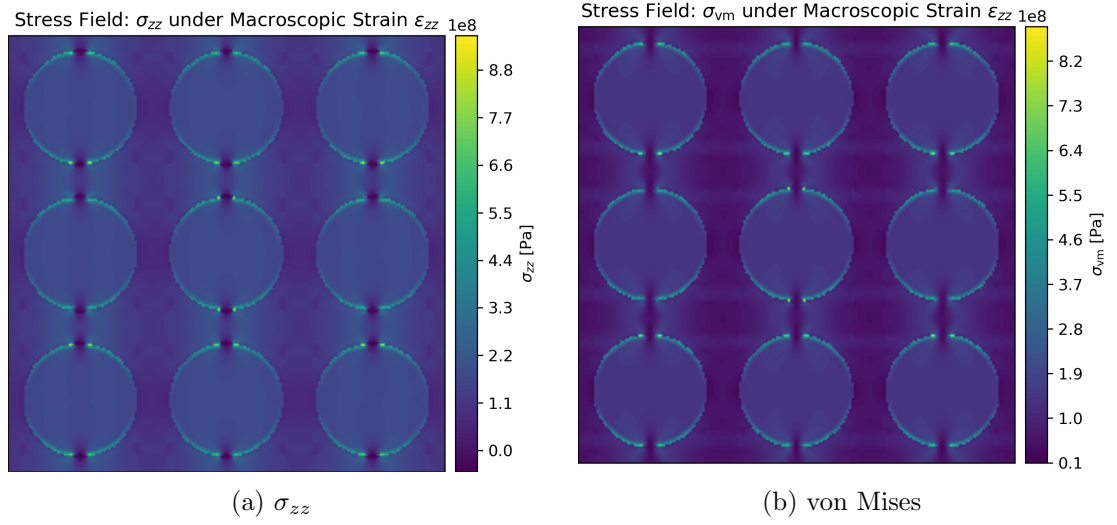
(b) von Mises

Figure 6.16: Stress fields at the mid-plane of the RVE under macroscopic strain ε_{zz}. (a) Stress component σ_{zz}. (b) von Mises stress σ_{vm}.

(a) σ_{yz}
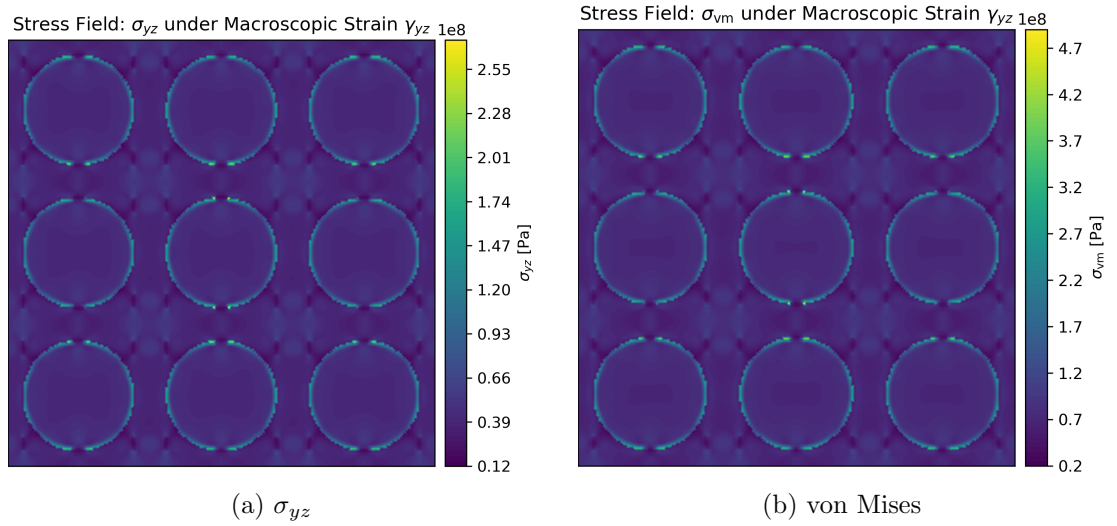
(b) von Mises

Figure 6.17: Stress fields at the mid-plane of the RVE under macroscopic strain γ_{yz}. (a) Stress component σ_{yz}. (b) von Mises stress σ_{vm}.

(a) σ_{zx} (b) von Mises

Figure 6.18: Stress fields at the mid-plane of the RVE under macroscopic strain γ_{zx}. (a) Stress component σ_{zx}. (b) von Mises stress σ_{vm}.

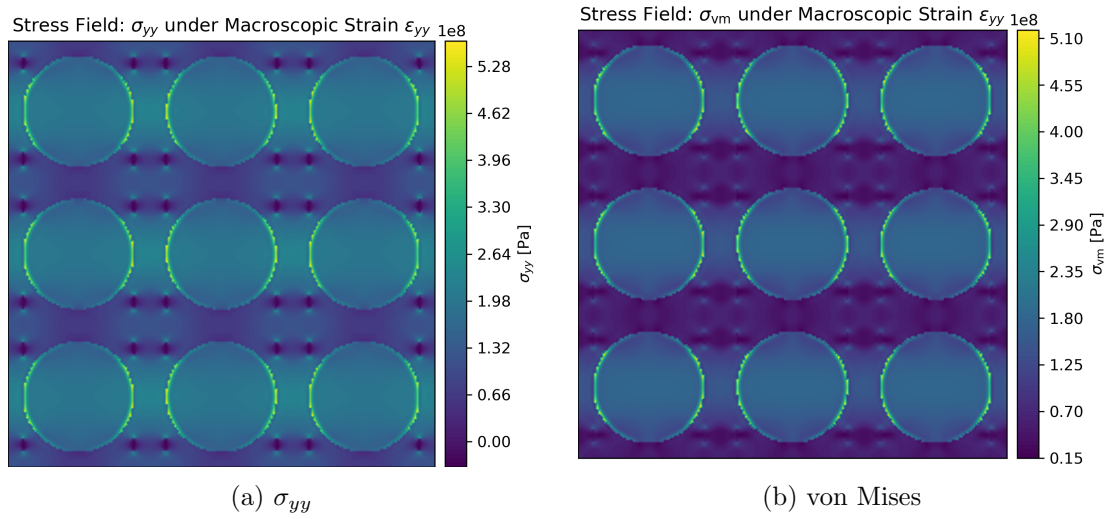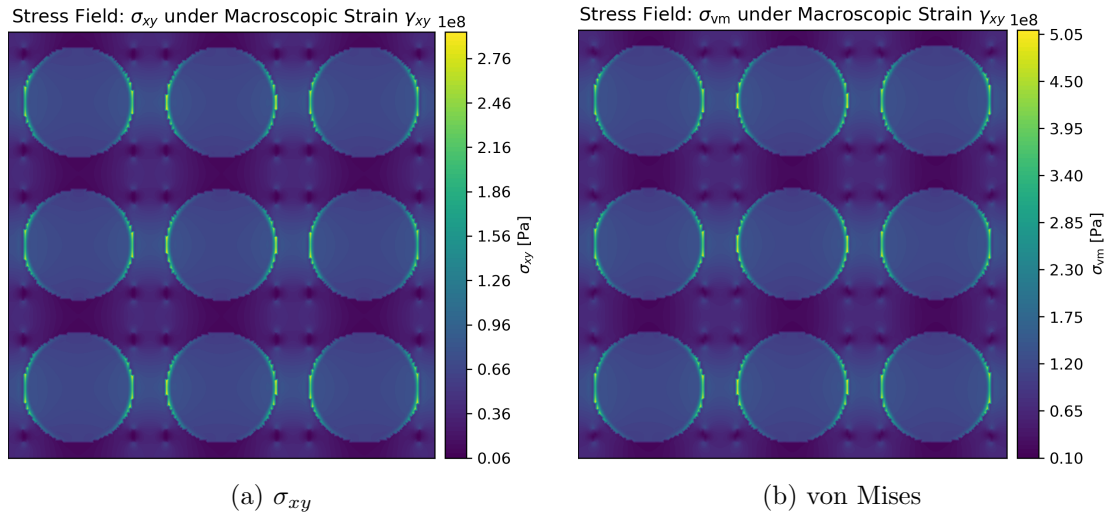(a) σ_{xy} (b) von Mises

Figure 6.19: Stress fields at the mid-plane of the RVE under macroscopic strain γ_{xy}. (a) Stress component σ_{xy}. (b) von Mises stress σ_{vm}.

As observed, the stress fields display strong periodicity consistent with the applied boundary conditions. Stress concentrations develop at fibre-matrix interfaces, especially around the debonding regions, and at crack tips where local stress gradients are amplified. In zones where multiple cracks are in close proximity, clear interaction effects emerge, leading to intensified stress bands. The von Mises stress fields further highlight regions of combined stress concentrations.

Notably, in the thickness direction (x-axis), stresses are predominantly carried by the boron fibres, as revealed in the σ_{xx}-dominated loading case. Surrounding matrix regions show comparatively lower stress levels, consistent with expected load transfer behaviour in fibre-reinforced composites. This distribution pattern, explicitly visualised due to the 3D formulation, reinforces the physical fidelity of the simulation and validates the extension from two-dimensional to fully three-dimensional modelling.

While no direct analytical comparison is included, the observed stress distributions and crack interaction effects align well with prior literature findings [49, 50] and the two-dimensional results presented in Chapter 5.

All plots are presented in physical units (Pa), without normalisation, preserving the absolute scale of stress magnitudes. This benchmark thus provides qualitative validation of local stress transfer and interaction phenomena in damaged composite microstructures under three-dimensional conditions.

Chapter Summary

This chapter extended the peridynamic computational homogenisation framework to fully three-dimensional RVEs, demonstrating its robustness and accuracy across a range of benchmark scenarios. Through systematic evaluation of effective stiffness tensors, interface treatment strategies, damage-induced property degradation, and local stress transfer phenomena, the method's applicability to complex 3D composite microstructures was established.

Benchmark V validated the framework's ability to predict anisotropic stiffness matrices for periodic fibre composites, with results aligning closely to theoretical expectations and confirming resolution convergence. Benchmark VI investigated effective bulk and shear moduli of glass-epoxy particulate composites across varying fibre volume fractions, showing good agreement with analytical models such as Mori-Tanaka and Hashin-Shtrikman bounds. Benchmark VII examined effective stiffness degradation in cracked matrix RVEs, capturing both progressive and abrupt stiffness loss trends consistent with analytical solutions and prior two-dimensional findings. Finally, Bench-

mark VIII visualised three-dimensional stress fields in a complex cracked composite, highlighting load transfer mechanisms, stress localisation, and crack interaction effects.

Collectively, the results presented in this chapter confirm the accuracy, flexibility, and physical consistency of the proposed three-dimensional peridynamic homogenisation approach, establishing a strong foundation for its application to multi-physics and damage modelling tasks in subsequent studies.

Chapter 7

Conclusion

This thesis presents a robust and scalable computational framework for the homogenisation of cracked composite materials using ordinary state-based peridynamics. By integrating raytracing-based bond-breaking, volumetric periodic boundary conditions, harmonic mean interface sampling, and strain-controlled homogenisation, the framework enables accurate analysis of complex microstructures in both two and three dimensions.

A principal novelty of this work is the generalised bond-breaking algorithm based on a modified Möller-Trumbore raytracing scheme, which transforms arbitrary cracks into triangle mesh surfaces, allowing natural and robust handling of complex damage morphologies without specialised treatments. The use of volumetric periodic boundary conditions ensures physically representative boundary behaviour for both 2D and 3D RVEs, while the MPI-based parallel implementation allows efficient simulation of large-scale problems. The framework exhibits near-linear computational scaling for problems exceeding approximately 1.3 million degrees of freedom, demonstrating its practical suitability for high-resolution 3D analyses.

Comprehensive benchmarks were conducted for both 2D and 3D composites, validating the framework against analytical solutions and established micromechanical results. Stress field visualisations, effective stiffness predictions, and crack interaction analyses confirm the method's robustness and physical fidelity across a wide range of scenarios.

Despite these strengths, the current framework is limited to purely elastic materials and neglects plasticity and temperature-dependent effects. Future developments could incorporate thermoelasticity, plasticity, progressive damage evolution, and multiscale coupling to extend the method's applicability to more complex material behaviours.

In summary, this thesis delivers a novel, general-purpose, and scalable computational approach for peridynamic homogenisation of cracked composites, combining advanced modelling features, validated physical fidelity, and parallel efficiency, while identifying clear directions for future enhancement.

Appendix A

Notation

Symbol	Description
$\Omega(t)$	Domain of the current (deformed) configuration at time t
Ω_0	Domain of the reference (undeformed) configuration
d	Spatial dimension, $d \in \{1, 2, 3\}$
$\mathbf{x}, \mathbf{x}'$	Material points in the reference configuration
dV'	Differential volume element associated with point $\mathbf{x}'$
$\rho(\mathbf{x})$	Mass density at material point $\mathbf{x}$
δ	Peridynamic horizon defining the extent of nonlocal interactions
$H_{\mathbf{x}}$	Family of material point $\mathbf{x}$ (i.e., points within horizon δ)
$\mathbf{u}(\mathbf{x})$	Displacement vector at point $\mathbf{x}$
$\ddot{\mathbf{u}}(\mathbf{x}, t)$	Acceleration vector at point $\mathbf{x}$ and time t
$\mathbf{y}, \mathbf{y}'$	Deformed positions of $\mathbf{x}$ and $\mathbf{x}'$
$\mathbf{y}' - \mathbf{y}$	Relative position vector in the deformed configuration
$\mathbf{x}' - \mathbf{x}$	Relative position vector in the reference configuration
S	Bond stretch (relative elongation of a bond)
$\underline{\mathbf{Y}}$	Deformation state operator (maps reference bond to deformed bond)
$\underline{\mathbf{T}}$	Force state operator (returns force density vector)
$\mathbf{t}, \mathbf{t}'$	Force density vectors exchanged between particles
$\mathbf{f}$	Symmetric pairwise force density (bond-based formulation)

Appendix A. Notation

Notation (continued)

Symbol	Description
$\mathbf{b}(\mathbf{x}, t)$	Body force density at point $\mathbf{x}$ and time t
C	Auxiliary parameter in bond-based peridynamics
$\overline{\mathbf{C}}, \overline{C}_{ijkl}$	Effective stiffness tensor
$\overline{\mathbf{S}}, \overline{S}_{ijkl}$	Effective compliance tensor
A, B	Auxiliary parameters in the ordinary state-based peridynamic model
W, W'	Strain energy densities at points $\mathbf{x}$ and $\mathbf{x}'$
θ, θ'	Dilatation at points $\mathbf{x}$ and $\mathbf{x}'$
a, b, d	Peridynamic material parameters calibrated to E, ν, and δ
E	Young's modulus
ν	Poisson's ratio
K	Bulk modulus
G	Shear modulus
λ, μ	Lamé parameters
$\boldsymbol{F}$	Deformation gradient
$\boldsymbol{\varepsilon}$	Strain tensor in Voigt notation
$\boldsymbol{\sigma}$	Cauchy stress tensor in Voigt notation
$\boldsymbol{D}$	Elastic stiffness matrix in Voigt notation
$\boldsymbol{D}^{ps}$	Plane stress stiffness matrix in Voigt notation
$\boldsymbol{D}^{pe}$	Plane strain stiffness matrix in Voigt notation
σ_{ij}	Components of the stress tensor
ε_{ij}	Components of the strain tensor
s	Applied shear or isotropic strain magnitude
h	Out-of-plane thickness (used in 2D plane stress problems)
r_s, θ_s, ϕ_s	Spherical coordinates for 3D numerical integration
r_p, θ_p	Polar coordinates for 2D numerical integration
ω, ω'	Influence functions (bond weighting functions)
V, V'	Volumes associated with points $\mathbf{x}$ and $\mathbf{x}'$

Appendix A. Notation

Notation (continued)

Symbol	Description
W^{ccm}_{shear}	Classical strain energy under simple shear
W^{pd}_{shear}	Peridynamic strain energy under simple shear
W^{ccm}_{normal}	Classical strain energy under isotropic expansion
W^{pd}_{normal}	Peridynamic strain energy under isotropic expansion
θ^{ccm}_{shear}	Classical dilatation under simple shear
θ^{ccm}_{normal}	Classical dilatation under isotropic expansion
θ^{pd}_{shear}	Peridynamic dilatation under simple shear
θ^{pd}_{normal}	Peridynamic dilatation under isotropic expansion
$s^{b}_{k,m}$	Correction factor for parameter b at particle k, with direction $m \in \{x, y, z\}$
:	Double contraction operator for tensors
$\overline{\boldsymbol{\varepsilon}}, \overline{\boldsymbol{\sigma}}$	Macroscopic strain and stress tensors
∇	Gradient operator
Δ	Laplace operator
∂	Partial derivative
$\mathbb{R}$	Set of real numbers
$\boldsymbol{I}$	Fourth-order identity tensor
$\boldsymbol{S}$	Fourth-order Eshelby tensor

Appendix B

Peridynamic Parameter Derivation under Plane Strain Conditions

In two-dimensional plane strain conditions, the stress and strain vectors in Voigt notation reduce to:

$$\boldsymbol{\sigma} = \begin{bmatrix} \sigma_{xx} & \sigma_{yy} & \sigma_{xy} \end{bmatrix}^T, \tag{B.1a}$$

$$\boldsymbol{\varepsilon} = \begin{bmatrix} \varepsilon_{xx} & \varepsilon_{yy} & \varepsilon_{xy} \end{bmatrix}^T. \tag{B.1b}$$

The corresponding stiffness matrix, obtained by inverting the compliance matrix, is:

$$\mathbf{D}^{pe} = \frac{E}{(1+\nu)(1-2\nu)} \begin{bmatrix} 1-\nu & \nu & 0 \\ \nu & 1-\nu & 0 \\ 0 & 0 & \frac{1-2\nu}{2} \end{bmatrix}. \tag{B.2}$$

To determine the peridynamic material parameters under plane strain, we follow the same calibration procedure as in subsection 2.4.2 by considering two deformation modes: simple shear and isotropic expansion. In each case, we compute the classical strain energy density and match it to the corresponding peridynamic expression.

Simple Shear Analysis

Appendix B. Peridynamic Parameter Derivation under Plane Strain Conditions

To determine the peridynamic parameter b, consider a simple shear strain state:

$$\varepsilon = \begin{bmatrix} 0 & 0 & s \end{bmatrix}^T. \tag{B.3}$$

This corresponds to the deformation gradient:

$$\boldsymbol{F} = \begin{bmatrix} 1 & s \\ 0 & 1 \end{bmatrix}. \tag{B.4}$$

The resulting stress vector is:

$$\boldsymbol{\sigma} = \mathbf{D}^{pe}\varepsilon = \begin{bmatrix} 0 & 0 & \frac{Es}{2(1+\nu)} \end{bmatrix}^T. \tag{B.5}$$

The strain energy density from classical continuum mechanics is:

$$W_{shear}^{ccm} = \frac{1}{2}\boldsymbol{\sigma}^T\varepsilon = \frac{Es^2}{4(1+\nu)}. \tag{B.6}$$

Since simple shear produces no volumetric change:

$$\theta_{shear}^{ccm} = 0. \tag{B.7}$$

We now evaluate the same deformation under the peridynamic formulation using polar coordinate (r_p, θ_p), with:

$$\left|\mathbf{x}' - \mathbf{x}\right| = r_p, \tag{B.8a}$$

$$\mathbf{x}' - \mathbf{x} = \begin{bmatrix} r_p\cos\theta_p \\ r_p\sin\theta_p \end{bmatrix}. \tag{B.8b}$$

The deformed bond under simple shear becomes:

$$\mathbf{y}' - \mathbf{y} = \boldsymbol{F}\cdot(\mathbf{x}' - \mathbf{x}) = \begin{bmatrix} r_p\cos\theta_p + sr_p\sin\theta_p \\ r_p\sin\theta_p \end{bmatrix}, \tag{B.9a}$$

$$\left|\mathbf{y}' - \mathbf{y}\right| \approx [1 + (\sin\theta_p\cos\theta_p)s]\,r_p. \tag{B.9b}$$

Appendix B. Peridynamic Parameter Derivation under Plane Strain Conditions

The bond length is linearised for small shear strains $s \ll 1$ using a first-order Taylor expansion, with higher-order terms $O(s^2)$ neglected.

The volume element in polar coordinate (r_p, θ_p) is:

$$dV' = r_p \, d\theta_p \, dr_p. \tag{B.10}$$

The corresponding peridynamic strain energy density in (2.13), expressed in integral form for simple shear is:

$$
\begin{aligned}
W_{shear}^{pd} &= b \int_{H_{\mathbf{x}}} \omega \left(\left| \mathbf{y}' - \mathbf{y} \right| - \left| \mathbf{x}' - \mathbf{x} \right| \right)^2 \, dV' \\
&= b \int_0^\delta \int_0^{2\pi} \frac{\delta}{r_p} \left\{ [1 + (\sin\theta_p \cos\theta_p)s] \, r_p - r_p \right\}^2 r_p \, d\theta_p \, dr_p.
\end{aligned}
\tag{B.11}
$$

Exploiting isotropy of the circular horizon, the angular integration yields a constant factor, independent of radius. The strain energy therefore becomes:

$$W_{shear}^{pd} = b \, \frac{\pi}{4} \, s^2 \int_0^\delta \delta \, r_s^2 \, dr_s, \tag{B.12}$$

integrating yields:

$$W_{shear}^{pd} = \frac{\pi \delta^4 s^2}{12} b. \tag{B.13}$$

Equating (B.13) with the classical strain energy density (B.6) gives:

$$b = \frac{3E}{\pi \delta^4 (1 + \nu)}. \tag{B.14}$$

Isotropic Expansion Analysis

To derive parameters a and d, consider the isotropic expansion:

$$\boldsymbol{\varepsilon} = \begin{bmatrix} s & s & 0 \end{bmatrix}^T. \tag{B.15}$$

110

Appendix B. Peridynamic Parameter Derivation under Plane Strain Conditions

This corresponds to the deformation gradient:

$$\boldsymbol{F} = \begin{bmatrix} 1+s & 0 \\ 0 & 1+s \end{bmatrix}. \tag{B.16}$$

The resulting stress vector is:

$$\boldsymbol{\sigma} = \mathbf{D}^{pe}\boldsymbol{\varepsilon} = \begin{bmatrix} \frac{Es}{(1+\nu)(1-2\nu)} \\ \frac{Es}{(1+\nu)(1-2\nu)} \\ 0 \end{bmatrix}. \tag{B.17}$$

The classical strain energy density and dilatation become:

$$W_{normal}^{ccm} = \frac{1}{2}\boldsymbol{\sigma}^{T}\boldsymbol{\varepsilon} = \frac{Es^2}{(1+\nu)(1-2\nu)}, \tag{B.18}$$

$$\theta_{normal}^{ccm} = 2s. \tag{B.19}$$

The deformed bond under isotropic expansion becomes:

$$\mathbf{y}' - \mathbf{y} = \boldsymbol{F} \cdot (\mathbf{x}' - \mathbf{x}) = \begin{bmatrix} r_p \cos\theta_p(1+s) \\ r_p \sin\theta_p(1+s) \end{bmatrix}, \tag{B.20a}$$

$$\left| \mathbf{y}' - \mathbf{y} \right| = (1+s)r_p. \tag{B.20b}$$

The volume element in polar coordinate (r_p, θ_p) is:

$$dV' = r_p \, d\theta_p \, dr_p. \tag{B.21}$$

The corresponding peridynamic strain energy density in (2.13), expressed in integral form for isotropic expansion is:

$$W_{normal}^{pd} = a(\theta_{normal}^{ccm})^2 + b \int_{H_{\mathbf{x}}} \omega \left(\left| \mathbf{y}' - \mathbf{y} \right| - \left| \mathbf{x}' - \mathbf{x} \right| \right)^2 dV'$$

$$= a(2s)^2 + b \int_0^{\delta} \int_0^{2\pi} \frac{\delta}{r_p} \left[(1+s)r_p - r_p \right]^2 r_p \, d\theta_p \, dr_p. \tag{B.22}$$

111

Appendix B. Peridynamic Parameter Derivation under Plane Strain Conditions

Since the integrand is independent of θ_p, the angular integration yields a factor of 2π. Integrating yields:

$$W^{pd}_{normal} = 4as^2 + \frac{2}{3}\pi b\delta^4 s^2. \tag{B.23}$$

The corresponding peridynamic dilatation in (2.14), expressed in integral form for isotropic expansion is:

$$
\begin{aligned}
\theta^{pd}_{normal} &= d \int_{H_\mathbf{x}} \omega\left(\left|\mathbf{y}' - \mathbf{y}\right| - \left|\mathbf{x}' - \mathbf{x}\right|\right)\left(\frac{\mathbf{y}' - \mathbf{y}}{|\mathbf{y}' - \mathbf{y}|} \cdot \frac{\mathbf{x}' - \mathbf{x}}{|\mathbf{x}' - \mathbf{x}|}\right) dV' \\
&= d \int_0^\delta \int_0^{2\pi} \frac{\delta}{r_p}\left[(1+s)r_p - r_p\right]\left(\frac{\boldsymbol{r_p}}{r_p} \cdot \frac{\boldsymbol{r_p}}{r_p}\right) r_p\, d\theta_p\, dr_p,
\end{aligned} \tag{B.24}
$$

Since $\mathbf{y}' - \mathbf{y}$ is collinear with $\mathbf{x}' - \mathbf{x}$ under isotropic expansion, the unit vectors are identical and $(\boldsymbol{r_p}/r_p) \cdot (\boldsymbol{r_p}/r_p) = 1$. The integrand is therefore independent of angular variables. Integrating yields:

$$\theta^{pd}_{normal} = \pi d\delta^3 s. \tag{B.25}$$

By equating (B.19) with (B.25), the expression for d is obtained:

$$d = \frac{2}{\pi\delta^3}. \tag{B.26}$$

Similarly, equating (B.18) with (B.23) yields the relationship between a and b. Solving these equations gives:

$$a = \frac{E(4\nu - 1)}{4(1+\nu)(1-2\nu)}. \tag{B.27}$$

Summary of Parameters in Plane Strain

$$a = \frac{E(4\nu - 1)}{4(1+\nu)(1-2\nu)}, \tag{B.28a}$$

$$b = \frac{3E}{\pi\delta^4(1+\nu)}, \tag{B.28b}$$

$$d = \frac{2}{\pi \delta^3}.$$

(B.28c)

$$d = \frac{2}{\pi \delta^3}.$$

(B.28c)

Bibliography

[1] S. A. Silling, "Reformulation of elasticity theory for discontinuities and long-range forces," *Journal of the Mechanics and Physics of Solids*, vol. 48, no. 1, pp. 175–209, 2000.

[2] S. A. Silling, M. Epton, O. Weckner, J. Xu, and E. Askari, "Peridynamic states and constitutive modeling," *Journal of Elasticity*, vol. 88, no. 2, pp. 151–184, 2007.

[3] E. Madenci, A. Barut, and N. Phan, "Peridynamic unit cell homogenization for thermoelastic properties of heterogenous microstructures with defects," *Composite Structures*, vol. 188, pp. 104–115, 2018.

[4] V. A. Buryachenko, *Local and nonlocal micromechanics of heterogeneous materials*. Springer, 2022.

[5] S. A. Silling and E. Askari, "A meshfree method based on the peridynamic model of solid mechanics," *Computers & structures*, vol. 83, no. 17-18, pp. 1526–1535, 2005.

[6] S. A. Silling and R. B. Lehoucq, "Peridynamic theory of solid mechanics," *Advances in applied mechanics*, vol. 44, pp. 73–168, 2010.

[7] E. Madenci and E. Oterkus, *Peridynamic Theory and Its Applications*. Springer Science & Business Media, 2013.

[8] S. A. Silling and R. B. Lehoucq, "Convergence of peridynamics to classical elasticity theory," *Journal of Elasticity*, vol. 93, pp. 13–37, 2008.

Bibliography

[9] S. A. Silling and F. Bobaru, "Peridynamic modeling of membranes and fibers," *International Journal of Non-Linear Mechanics*, vol. 40, no. 2-3, pp. 395–409, 2005.

[10] A. Candaş, E. Oterkus, and C. E. İmrak, "Peridynamic simulation of dynamic fracture in functionally graded materials subjected to impact load," *Engineering with Computers*, vol. 39, no. 1, pp. 253–267, 2023.

[11] A. Candaş, E. Oterkus, and C. E. Imrak, "Ordinary state-based peridynamic modelling of crack propagation in functionally graded materials with micro cracks under impact loading," *Mechanics of Advanced Materials and Structures*, vol. 31, no. 30, pp. 13 502–13 517, 2024.

[12] Y. K. Galadima, W. Xia, E. Oterkus, and S. Oterkus, "A computational homogenization framework for non-ordinary state-based peridynamics," *Engineering with Computers*, vol. 39, no. 1, pp. 461–487, 2023.

[13] W. Xia, E. Oterkus, and S. Oterkus, "Ordinary state-based peridynamic homogenization of periodic micro-structured materials," *Theoretical and Applied Fracture Mechanics*, vol. 113, p. 102960, 2021.

[14] R. W. Macek and S. A. Silling, "Peridynamics via finite element analysis," *Finite elements in analysis and design*, vol. 43, no. 15, pp. 1169–1178, 2007.

[15] V. Lecomte, "Handbook of geometry for competitive programmers," *Draft, October*, vol. 14, p. 2018, 2018.

[16] T. Möller and B. Trumbore, "Fast, minimum storage ray-triangle intersection," *Journal of Graphics Tools*, vol. 2, no. 1, pp. 21–28, 1997.

[17] W. Xia, E. Oterkus, and S. Oterkus, "3-dimensional bond-based peridynamic representative volume element homogenisation," *Physical Mesomechanics*, vol. 24, pp. 541–547, 2021.

[18] R. Hill, "Elastic properties of reinforced solids: some theoretical principles," *Journal of the Mechanics and Physics of Solids*, vol. 11, no. 5, pp. 357–372, 1963.

[19] ——, "On constitutive macro-variables for heterogeneous solids at finite strain," *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, vol. 326, no. 1565, pp. 131–147, 1972.

[20] S. Nemat-Nasser and M. Hori, *Micromechanics: overall properties of heterogeneous materials.* Elsevier, 1993, vol. 37.

[21] P. Suquet, "Elements of homogenization for inelastic solid mechanics," *Homogenization techniques for composite media*, vol. 272, pp. 193–278, 1987.

[22] J. Fish, K. Shek, M. Pandheeradi, and M. S. Shephard, "Computational plasticity for composite structures based on mathematical homogenization: Theory and practice," *Computer methods in applied mechanics and engineering*, vol. 148, no. 1-2, pp. 53–73, 1997.

[23] M. G. Geers, V. G. Kouznetsova, and W. Brekelmans, "Multi-scale computational homogenization: Trends and challenges," *Journal of computational and applied mathematics*, vol. 234, no. 7, pp. 2175–2182, 2010.

[24] C. Miehe, J. Schröder, and J. Schotte, "Computational homogenization analysis in finite plasticity simulation of texture development in polycrystalline materials," *Computer methods in applied mechanics and engineering*, vol. 171, no. 3-4, pp. 387–418, 1999.

[25] C. Miehe and A. Koch, "Computational micro-to-macro transitions of discretized microstructures undergoing small strains," *Archive of Applied Mechanics*, vol. 72, pp. 300–317, 2002.

[26] C. Miehe, "Computational micro-to-macro transitions for discretized microstructures of heterogeneous materials at finite strains based on the minimization of averaged incremental energy," *Computer methods in applied mechanics and engineering*, vol. 192, no. 5-6, pp. 559–591, 2003.

[27] W. Voigt, "Ueber die beziehung zwischen den beiden elasticitätsconstanten isotroper körper," *Annalen der physik*, vol. 274, no. 12, pp. 573–587, 1889.

Bibliography

[28] A. Reuß, "Berechnung der fließgrenze von mischkristallen auf grund der plastizitätsbedingung für einkristalle." *ZAMM-Journal of Applied Mathematics and Mechanics/Zeitschrift für Angewandte Mathematik und Mechanik*, vol. 9, no. 1, pp. 49–58, 1929.

[29] Z. Hashin and S. Shtrikman, "A variational approach to the theory of the elastic behaviour of multiphase materials," *Journal of the Mechanics and Physics of Solids*, vol. 11, no. 2, pp. 127–140, 1963.

[30] T. Mori and K. Tanaka, "Average stress in matrix and average elastic energy of materials with misfitting inclusions," *Acta metallurgica*, vol. 21, no. 5, pp. 571–574, 1973.

[31] Y. Benveniste, "A new approach to the application of mori-tanaka's theory in composite materials," *Mechanics of materials*, vol. 6, no. 2, pp. 147–157, 1987.

[32] J. D. Eshelby, "The determination of the elastic field of an ellipsoidal inclusion, and related problems," *Proceedings of the royal society of London. Series A. Mathematical and physical sciences*, vol. 241, no. 1226, pp. 376–396, 1957.

[33] R. Hill, "A self-consistent mechanics of composite materials," *Journal of the Mechanics and Physics of Solids*, vol. 13, no. 4, pp. 213–222, 1965.

[34] R. Christensen and K. Lo, "Solutions for effective shear properties in three phase sphere and cylinder models," *Journal of the Mechanics and Physics of Solids*, vol. 27, no. 4, pp. 315–330, 1979.

[35] G. P. Tandon and G. J. Weng, "The effect of aspect ratio of inclusions on the elastic properties of unidirectionally aligned composites," *Polymer composites*, vol. 5, no. 4, pp. 327–333, 1984.

[36] J. H. Affdl and J. Kardos, "The halpin-tsai equations: a review," *Polymer Engineering & Science*, vol. 16, no. 5, pp. 344–352, 1976.

Bibliography

[37] R. B. Lehoucq and S. A. Silling, "Force flux and the peridynamic stress tensor," *Journal of the Mechanics and Physics of Solids*, vol. 56, no. 4, pp. 1566–1577, 2008.

[38] E. Gabriel, G. E. Fagg, G. Bosilca, T. Angskun, J. J. Dongarra, J. M. Squyres, V. Sahay, P. Kambadur, B. Barrett, A. Lumsdaine *et al.*, "Open mpi: Goals, concept, and design of a next generation mpi implementation," in *European Parallel Virtual Machine/Message Passing Interface Users' Group Meeting*. Springer, 2004, pp. 97–104.

[39] S. Balay, S. Abhyankar, M. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. Gropp *et al.*, *PETSc users manual*. Argonne National Laboratory, 2019.

[40] G. L. Sleijpen and H. A. Van der Vorst, "An overview of approaches for the stable computation of hybrid bicg methods," *Applied numerical mathematics*, vol. 19, no. 3, pp. 235–254, 1995.

[41] G. L. Sleijpen, H. A. Van der Vorst, and D. R. Fokkema, "Bicgstab (l) and other hybrid bi-cg methods," *Numerical Algorithms*, vol. 7, no. 1, pp. 75–109, 1994.

[42] D. R. Fokkema, *Enhanced implementation of BiCGstab (l) for solving linear systems of equations*. Universiteit Utrecht. Mathematisch Instituut, 1996.

[43] C.-T. Sun and R. S. Vaidya, "Prediction of composite properties from a representative volume element," *Composites science and Technology*, vol. 56, no. 2, pp. 171–179, 1996.

[44] W. Yu and T. Tang, "Variational asymptotic method for unit cell homogenization of periodically heterogeneous materials," *International Journal of Solids and Structures*, vol. 44, no. 11-12, pp. 3738–3755, 2007.

[45] J. Aboudi, M.-J. Pindera, and S. Arnold, "Linear thermoelastic higher-order theory for periodic multiphase materials," *J. Appl. Mech.*, vol. 68, no. 5, pp. 697–707, 2001.

Bibliography

[46] C. Chamis, "Simplified composite micromechanics equations of hygral, thermal, and mechanical properties," *SAMPE Quarterly*, vol. 15, no. ISSN: 0036-0821, 1984.

[47] D. Kenaga, J. F. Doyle, and C. Sun, "The characterization of boron/aluminum composite in the nonlinear range as an orthotropic elastic-plastic material," *Journal of Composite Materials*, vol. 21, no. 6, pp. 516–531, 1987.

[48] X. Markenscoff and C. Dascalu, "Asymptotic homogenization analysis for damage amplification due to singular interaction of micro-cracks," *Journal of the Mechanics and Physics of Solids*, vol. 60, no. 8, pp. 1478–1485, 2012.

[49] W. Zhu and Q. Dong, "Effective elastic properties of cracked composites with periodically distributed particulates," *Mechanics of Advanced Materials and Structures*, vol. 29, no. 16, pp. 2267–2277, 2022.

[50] K. Zhou, "Elastic field and effective moduli of periodic composites with arbitrary inhomogeneity distribution," *Acta Mechanica*, vol. 223, no. 2, pp. 293–308, 2012.

Bibliography